

THE UNIVERSITY OF CHICAGO

MEASUREMENT OF THE CMB POLARIZATION AT 95 GHZ FROM QUIET

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
IMMANUEL DAVID BUDER

CHICAGO, ILLINOIS

AUGUST 2012

Copyright © 2012 by Immanuel David Buder

All rights reserved

to Mr. Wallin, who never doubted

and

to Bruce, who made it possible

Table of Contents

List of Figures	viii
List of Tables	x
Acknowledgments	xi
Abstract	xiii
1 Introduction	1
1.1 Inflationary Cosmology	1
1.2 The Cosmic Microwave Background	6
1.2.1 CMB Measurements and Their Implications	7
1.2.2 CMB Polarization Measurements	9
1.3 The QUIET Experiment	12
1.3.1 Module Operation Principles	15
1.3.2 Optics	18
1.3.3 Modules	22
1.3.4 Cryostat	25
2 Electronics	29
2.1 Passive Interfaces	29
2.2 Bias	36
2.2.1 Temperature Regulation	41
2.2.2 WPID Output Monitoring Board	44
2.3 Readout	44
3 Data Management	51
3.1 Online Software	52
3.1.1 <code>adc_server</code>	53
3.1.2 <code>DataCompilation</code>	59
3.1.3 <code>ReceiverControl</code>	65
3.1.4 <code>DataArc</code>	66
3.2 Timing Alignment	66
3.3 Control Software	67
3.3.1 <code>Receiver Control Panel</code>	67
3.3.2 Schedule Making	68
3.3.3 Real-time Monitoring	69
3.4 Data Transfer and Storage	69
3.5 Database	71

4 Observations	72
4.1 Site	72
4.2 Regular Observations	74
4.2.1 CMB Observations	74
4.2.2 Calibration Observations	77
4.3 Season-end Calibrations	81
4.4 Noteworthy Events	81
4.5 Routine Checks	82
4.5.1 Enclosure Temperature	83
4.5.2 Data Management Status	83
5 Analysis	88
5.1 Calibration	89
5.1.1 Beam Shape	89
5.1.2 Responsivity	91
5.1.3 Detector Angle	93
5.1.4 Pointing	94
5.1.5 Type-B-glitch Correction	97
5.2 Noise Modeling	101
5.2.1 Correlated 1/f Noise Investigation	103
5.2.2 Model of Noise Correlations	111
5.3 Time-ordered-data Processing	112
5.3.1 Filtering	114
5.3.2 Data Selection	114
5.4 Power-spectrum Calculation	117
5.4.1 Map Making	121
5.4.2 Pseudo- C_ℓ Calculation	121
5.4.3 Cross-correlation	123
5.5 Simulations	124
5.6 Validation Tests	125
5.6.1 Blind-analysis Strategy	127
5.6.2 Null Suite	128
5.6.3 Patch Comparison	134
5.6.4 Configuration Comparison	136
5.7 Systematic Errors	138
5.7.1 Miscalibration	139
5.7.2 Instrumental Polarization	142
5.7.3 Type-B-glitch Miscorrection	143
5.7.4 Sun Pickup	149
5.7.5 Type-O Glitching	150
5.7.6 Negligible Effects	151

6 Results and Discussion	153
6.1 Confidence Intervals	153
6.2 Comparison with Alternate Analysis Pipeline	155
6.3 Consistency with Λ CDM	155
6.4 Limit on Inflation	160
6.5 Conclusions	160
A DataCompilation Status Flags	163
B Online Software Details	170
B.1 <code>adc_server</code> Packet Status Bytes	170
B.2 <code>adc_sever</code> bias Thread Actions	171
B.3 <code>data_serve</code> Requests	173
B.4 DataCompilation–DataArc Communication	175
B.5 <code>ReceiverControl</code> Commands	176
C Timing Alignment Algorithm	180
D Receiver Database	182
D.1 Table ADC_MAPPING	182
D.2 Table BIASES	183
D.3 Table HOUSEKEEPING	183
D.4 Table mab_input	184
E Database and Data-quality Statistics	186
E.1 Database quiet_cal	186
E.1.1 Table w91_gain_model_each_runs	186
E.2 Database quiet_data	187
E.2.1 Table w91_ces	187
E.2.2 Table w91_l0_ces_intersection	187
E.2.3 Table w91_level0	188
E.2.4 Table w91_level1	189
E.2.5 Table w91_run_log	190
E.3 Database quiet_quality	191
E.3.1 Table w91_atm	191
E.3.2 Table w91_ces_usable	192
E.3.3 Table w91_diode_usable	193
E.3.4 Table w91_diodes	193
E.3.5 Table w91_full_correlation	194
E.3.6 Table w91_housekeeping	194
E.3.7 Table w91_scan	195
E.3.8 Table w91_software	197
E.3.9 Table w91_sun_spike	200
E.3.10 Table w91_timestream	200

E.3.11 Table w91_typeb	203
F Array Sensitivity Definition and Calculation	205
F.1 Diode Sensitivity	205
F.2 Diode-pair Sensitivity	205
F.3 Module Sensitivity	206
F.4 Array Sensitivity	207
G Null-test Definitions and Correlations	208
H List of Analysis Configurations	213
I Statistical Uncertainty Scaling Between Configurations	220
I.1 Configuration-difference Error Bar	220
I.2 Impact of CMB-instrument-noise Accidental Correlation	222
I.2.1 CMB-instrument-noise Correlation Toy Model	223
I.2.2 Multiple Sky Modes	224
I.2.3 EB Power Spectrum	225
J Type-B-glitch—correction-parameter File Format	227
References	230

List of Figures

1.1	Creation of Linear Polarization in the CMB	10
1.2	QUIET Instrument Deployed	14
1.3	QUIET Module Diagram	15
1.4	Field of View	19
1.5	Sidelobe Diagram	21
1.6	Module Photograph	24
1.7	Cryostat Internals	27
1.8	Cryostat Temperature Throughout the Season	28
2.1	Electronics Systems	30
2.2	Module Assembly Boards Tiling the Focal Plane	31
2.3	Array Interface Boards	33
2.4	AIB Cable Boards	34
2.5	Array Interface Board Weatherproofing	35
2.6	Phase-switch Transmission Curve	37
2.7	Phase-switch Modulated Data	38
2.8	Bias-Board Backplane	41
2.9	Electronics Enclosure	42
2.10	Electronics Enclosure Temperature Regulation System	43
2.11	WPID Output Monitoring Board Schematic	45
2.12	Demodulation Cartoon	47
2.13	Type-B-glitch Cartoon	48
2.14	Auxiliary Timing Board Schematic	50
3.1	Data Management System Overview	52
3.2	<code>adc_server</code> Radiometer Data Packets	56
3.3	Live:QUIET Automatic Photograph	70
4.1	Precipitable Water Vapor near the Site	73
4.2	Sky Fields	76
4.3	Scan Strategy Cartoon	78
4.4	Scan Trajectory	79
4.5	Electronics Enclosure Temperature During the Observing Season	84
4.6	Electronics Enclosure Temperature Stability	85
5.1	Beam Profile	90
5.2	1/f Correlation Matrix Eigenvalues	105
5.3	1/f Correlation Matrix Eigenvector and Demonstration of two 1/f Noise Components	106
5.4	Simulated 1/f Correlation Matrix Eigenvalues	107
5.5	1/f Noise Model χ^2	109
5.6	1/f Correlation Matrix Eigenvalues in Each Module	112
5.7	1/f Correlation Matrix Eigenvectors	113

5.8	RQ57Q1 TP with Two States	118
5.9	RQ57Q1 Noise-model Fit χ^2	118
5.10	DD vs. TP	119
5.11	Noise-spectrum Spikes	119
5.12	CES-diode with Noise Level Changes	120
5.13	Transfer Function	126
5.14	Pipeline Fiducial Simulations	127
5.15	Null-test Correlation Matrix	131
5.16	Distribution of χ_b^{null}	133
5.17	Excess Noise in Power Spectra Including Auto-correlations	135
5.18	Patch Comparison χ_p^2 Distribution	137
5.19	Power Spectrum Difference Between Different Data-selection Thresholds	138
5.20	Systematic Errors Summary	140
5.21	Type-B-glitch—misCorrection Instrumental Polarization	145
5.22	Type-B-misCorrection Systematic Error vs. Threshold	147
5.23	800-kHz-data Distribution	149
6.1	Polarization Power Spectra	156
6.2	Polarization Power Spectra of Each Patch	157
6.3	Window Functions	158
6.4	Power Spectra from Two Pipelines	159
6.5	Power Spectra Compared to Previous Experiments	161

List of Tables

2.1	Bias-board Locations	40
3.1	Summary of <code>adc_server</code> Threads	54
3.2	Summary of <code>DataCompilation</code> Threads	60
3.3	Summary of <code>DataCompilation</code> <code>data_server</code> Requests	62
3.4	Summary of <code>DataCompilation</code> and <code>DataArc</code> Communication	65
4.1	Observation Summary	75
4.2	CMB Patch Centers	77
4.3	Regular Calibration Observations	80
4.4	Electronics Enclosure Temperature Regulation Setpoints	83
5.1	Absolute-responsivity–model Time Divisions	92
5.2	Pointing-model Parameters	96
5.3	Data-selection Efficiency	120
5.4	Null-suite Probabilities to Exceed	128
B.1	<code>ReceiverControl</code> Commands	176
D.1	MAB–bias-board Mapping	185
D.2	Modules and ADC Board for Each MAB	185
G.1	Null-test Definitions	208

Acknowledgments

Here friend, have this!

Okay, thanks.

Team Potato

I could never have written this thesis without the support of many people. Here I thank as many as memory will permit.

First thanks go to my adviser Bruce Winstein, who was also the PI of QUIET. Without him I never would have gotten involved in CMB physics. Bruce taught me how to be a really careful scientist. He was also incredibly generous with his time and resourceful at getting me whatever I needed to get that science done. My greatest regret is that he never got to see this work finished.

My adviser Stephan Meyer stepped in at a time when we were all reeling from Bruce's death. Steve helped me understand all the things I thought I understood—but didn't. Steve was also really good at keeping me focused on the next steps at every stage in this process.

Akito Kusaka took on responsibility beyond what is normally expected of a post-doc. He kindly pointed out mistakes in every memo I sent him. Any ones remaining here are my own.

My thesis committee, Paolo Privitera, Michael Turner, and David Biron, provided many insightful comments and suggestions. This thesis is much better as a result.

Thanks to the rest of the Chicago QUIET group, Colin Bischoff, Alison Brizius, Yuji Chinone, Dan Kapner, and Osamu Tajima, for keeping things fun. My Fermilab colleagues, Hogan Nguyen and Donna Kubik, kept up the enthusiasm when the Chicago team was burnt out.

Thanks to the Q-band deployment team for putting up with me even when I was too low on oxygen to make any sense. Special thanks to the engineers and technician who

maintained and operated the telescope, José Cortés, Cristobal Jara, Freddy Muñoz, and Carlos Verdugo. Without them, we never would have gotten past the first mount stall. Thanks also to the other Atacama experiments, especially APEX, ACT, and ALMA, for supporting our operations in divers ways.

Thanks to everyone else in the QUIET Collaboration for numerous comments, discussions, and a whole lot of science.

My office mates, Cora Dvorkin, Sam Leitner, Denis Erkal, and Yin Li, provided job-seeking support, laughs, and daily encouragement.

Thanks to my classmates and friends in Chicago, who were uniformly, point-wise, and L2 supportive. Thanks to my Tpot friends, whose support was neither closed nor bounded.

My family thought I was doing something worthwhile, even when they didn't understand it.

Some work was performed on the Joint Fermilab-KICP Supercomputing Cluster, supported by grants from Fermilab, the Kavli Institute for Cosmological Physics, and the University of Chicago. Some work was performed on the Central Computing System, owned and operated by the Computing Research Center at KEK. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Some of the results in this paper have been derived using the HEALPix ([Górski et al., 2005](#)) package.

Last and least, I thank the mount, which lasted much longer than anyone thought it would.

Abstract

Despite the great success of precision cosmology, cosmologists cannot fully explain the initial conditions of the Universe. Inflation, an exponential expansion in the first $\approx 10^{-36}$ s, is a promising potential explanation. A generic prediction of inflation is odd-parity (B-mode) polarization in the cosmic microwave background (CMB). The Q/U Imaging Experiment (QUIET) aimed to limit or detect this polarization.

We built a coherent pseudo-correlation microwave polarimeter. An array of mass-produced “modules” populated the focal plane of a 1.4-m telescope. Each module had a sensitivity to polarization of $756 \mu\text{K}\sqrt{\text{s}}$ with a bandwidth of $10.7 \pm 1.1 \text{ GHz}$ centered at $94.5 \pm 0.8 \text{ GHz}$; the combined sensitivity was $87 \pm 7 \mu\text{K}\sqrt{\text{s}}$. We incorporated “deck” rotation, an absorbing ground screen, a new time-stream “double-demodulation” technique, and optimized optics into the design to reduce instrumental polarization. We observed with this instrument at the Atacama Plateau in Chile between August 2009 and December 2010. We collected 5336.9 hours of CMB observation and 1090 hours of astronomical calibration.

This thesis describes the analysis and results of these data. We characterized the instrument using the astronomical calibration data as well as purpose-built artificial sources. We developed noise modeling, filtering, and data selection following a blind-analysis strategy. Central to this strategy was a suite of 32 null tests, each motivated by a possible instrumental problem or systematic effect. We also evaluated the systematic errors in the blind stage of the analysis before the result was known. We then calculated the CMB power spectra using a pseudo- C_ℓ cross-correlation technique that suppressed contamination and made the result insensitive to noise bias. We measured the first three peaks of the E-mode spectrum at high significance and limited B-mode polarization ($r = 1.48^{+1.06}_{-0.92}$ at 68% confidence and $r < 3.3$ at 95% confidence). Systematic errors were well below ($r < 0.01$) our B-mode polarization limit. This systematic-error reduction was a strong demonstration of technology for application in more sensitive, next-generation CMB experiments.

Chapter 1

Introduction

In the beginning the Universe was created. This has made a lot of people very angry and been widely regarded as a bad move.

Douglas Adams

Q/U Imaging ExperimenT (QUIET) measured the cosmic microwave background (CMB) polarization with the goal of constraining a signal from inflation. The standard model of cosmology begins with inflation, an accelerated expansion of the Universe in the first $\approx 10^{-36}$ s that sets its initial conditions. Inflation generated a stochastic background of gravitational waves. These gravitational waves created a small polarization anisotropy in the CMB, the radiation released when the early Universe became un-ionized. We¹ measured the CMB polarization and isolated the signal potentially bearing evidence of inflation. This thesis describes the QUIET instrument, observations, data analysis, and results including a limit ($r < 3.3$) on the inflation signal.

1.1 Inflationary Cosmology

Inflation provides the initial conditions for the standard model of cosmology (Λ CDM, see e.g. Ryden (2002)). In the standard model, the Universe is assumed to be homogeneous and isotropic on large scales ($\gtrsim 100$ Mpc). In General Relativity, these conditions require a metric of the form

$$ds^2 = -c^2 dt^2 + a(t)^2 (dx^2 + dy^2 + dz^2). \quad (1.1)$$

1. Herein “we” means the members of the QUIET Collaboration, <http://quiet.uchicago.edu/>. “I” indicates particular contributions of the author.

This form is only valid for a nearly flat Universe, one of the assumptions of the standard model. The Universe expanded from a hot, dense initial state; thus $a(t)$ ² increases with time. In addition to the background metric, the standard model specifies the energy content and how it behaves. The evolution is governed by

$$\left(\frac{1}{a} \frac{da}{dt}\right)^2 = \frac{8\pi G}{3c^2} \rho \quad (1.2)$$

$$\frac{1}{a} \frac{d^2a}{dt^2} = -\frac{4\pi G}{3c^2} (\rho + 3p), \quad (1.3)$$

where ρ is the energy density and p is the pressure.

Inflation was originally proposed by [Guth \(1981\)](#) to explain puzzling observations (see also §1.2): the Universe is homogeneous even across regions that were causally disconnected. The Universe is (almost) flat, but a nearly flat Universe is an unstable solution. Guth's final motivation was the absence of magnetic monopoles, which are generally expected in grand unified theories. The generation of initial density perturbations is an additional motivation for an inflationary origin.

Although there are many models for the details of inflation, a simple, single-field, slow-roll model³ demonstrates the general features important for this thesis. Suppose there is a scalar field ϕ with mass m ⁴ and potential

$$V(\phi) = m^2 \phi^2 / 2. \quad (1.4)$$

The corresponding equations of motion are

$$\ddot{\phi} + 3H\dot{\phi} = -m^2\phi \quad (1.5)$$

2. Often called the “scale factor.”

3. This discussion follows [Linde \(2008\)](#).

4. In this example I chose units $M_p^{-2} = 8\pi G = 1$.

and

$$H^2 + \frac{k}{a^2} = (\dot{\phi}^2 + m^2\phi^2)/6, \quad (1.6)$$

where k is -1 , 0 , or 1 for an open, flat, or closed Universe, respectively. An overdot denotes a derivative with respect to time coordinate t , and $H \equiv \dot{a}/a$. If the initial value of ϕ is $\gg 1$ then the evolution is in the “slow-roll” regime with

$$\ddot{\phi} \ll 3H\dot{\phi}, \quad (1.7)$$

$$H^2 \gg k/a^2, \quad (1.8)$$

and

$$\dot{\phi}^2 \ll m^2\phi^2. \quad (1.9)$$

Eq. 1.8 implies that spatial curvature becomes negligible. In this regime the equations of motion simplify to

$$H = m\phi/\sqrt{6} \quad (1.10)$$

and

$$\dot{\phi} = -m\sqrt{2/3}. \quad (1.11)$$

In the slow-roll regime, the change of ϕ is negligible so

$$a(t) = e^{Ht} = e^{m\phi t/\sqrt{6}}. \quad (1.12)$$

This exponential expansion continues until $\phi \approx 1$. If the initial scale factor and field values are a_0 and ϕ_0 , respectively, then the total expansion caused by inflation is⁵

$$a/a_0 = e^{\phi_0^2/4}. \quad (1.13)$$

5. See §1.7 of [Linde \(2005\)](#) for a detailed derivation.

This model can explain the “puzzling observations” above. For realistic⁶ $m \approx 10^{-6}$, $a/a_0 \approx 10^{10^{10}}$. This expansion occurred in $\approx 10^{-36}$ s. Under a “worst case” assumption that the Universe before inflation had radius of curvature approximately the Planck length (10^{-33} cm), after inflation it would grow to $\approx 10^{10^{10}}$ cm. This is much larger than the present observable Universe, 10^{28} cm. Thus all observations cosmologists have made correspond to a small patch that was a tiny fraction of the total volume of the initial Universe and easily causally connected. The initial curvature is now unobservable because the present radius of curvature is much larger than any scale we can observe. For similar reasons, the expansion erases any evidence of initial inhomogeneity and so reduces the density of monopoles that it is unlikely to find one in our Hubble volume. At the end of inflation, ϕ decays into other elementary particles which interact and eventually reach thermal equilibrium.

Due to quantum fluctuations there were perturbations in the scalar field

$$\delta\phi \approx \frac{H}{2\pi}. \quad (1.14)$$

Differences in the value of ϕ caused inflation to end slightly earlier or later than the average value. The change in the time of the end of inflation is

$$\delta t = \delta\phi/\dot{\phi} \approx \frac{H}{2\pi\dot{\phi}}. \quad (1.15)$$

From Eq. 1.2,

$$\rho \propto H^2. \quad (1.16)$$

After the end of inflation, the Universe was radiation-dominated so

$$H \propto t^{-1}. \quad (1.17)$$

6. As shown below, the perturbation amplitude $\delta\rho/\rho$ depends on m . To be realistic, m must be consistent with the observed perturbation amplitude.

Combining Eqs. 1.14–1.17,

$$\delta H/H \approx \delta t/t \approx \delta \rho/\rho \approx \frac{H^2}{2\pi\dot{\phi}}. \quad (1.18)$$

The perturbation amplitude for a given wavelength depends on the fluctuations occurring when that wavelength was comparable to the horizon size. Since $\frac{H^2}{2\pi\dot{\phi}}$ changed slowly during the slow-roll regime, the spectrum of initial perturbations was nearly scale-invariant. Thus the combination of inflation and quantum fluctuations naturally produced the small anisotropy we observe today.

In addition to the density perturbations, inflation will in general create tensor perturbations (gravitational waves). I adopted the convention (Komatsu et al., 2011)

$$\Delta_R^2(k) \equiv \frac{k^3 \langle |R_k|^2 \rangle}{2\pi^2} = \Delta_R^2(k_0) (k/k_0)^{n_s-1} \quad (1.19)$$

$$\Delta_h^2(k) \equiv \frac{4k^3 \langle |h_k|^2 \rangle}{2\pi^2} = \Delta_h^2(k_0) (k/k_0)^{n_t} \quad (1.20)$$

for nearly scale-invariant⁷ spectra, where R_k is a perturbation of the scalar curvature with comoving wavenumber k , k_0 is a pivot scale⁸, n_s is the scalar spectral index, h_k is a tensor perturbation, and n_t is the tensor spectral index.

There is no observational evidence for tensor perturbations; however, many inflationary models predict them. The tensor-to-scalar ratio

$$r \equiv \frac{\Delta_h^2(k_0)}{\Delta_R^2(k_0)} \quad (1.21)$$

is a figure of merit for comparing tensor-perturbation searches and comparing experimental

7. $n_s \approx 1$ and $n_t \ll 1$.

8. In this thesis, $k_0 = 0.002 \text{ Mpc}^{-1}$.

results to different inflationary models. In the slow-roll regime,

$$r = 16\epsilon \equiv 8 \left(\frac{1}{V} \frac{dV}{d\phi} \right)^2. \quad (1.22)$$

In the $V = m^2\phi^2/2$ model, $r = 32/\phi^2$. Using Eq. 1.13, $r = 8/N$, where N is the number of e-folds of inflation for scales corresponding to the size of the observable Universe. For typical $N \approx 60$, $r \approx 0.13$. Regardless of the model, detection of non-zero r would be new, strong evidence for inflation. Moreover, a measurement of r would give information about the energy scale of inflation (Liddle & Lyth, 1993; Baumann et al., 2009)

$$V^{1/4} = 1.06 \times 10^{16} \text{ GeV} (r/0.01)^{1/4}. \quad (1.23)$$

The best existing limit is $r < 0.2$ (Komatsu et al., 2011). For reasons that will be explained below (§1.2.2), only $r \gtrsim 0.001$ is detectable in CMB polarization. Therefore, a measurement of r would indicate that inflation occurred at an energy comparable to the grand unification scale of particle physics. Although models with arbitrarily small r exist, current evidence and naturalness considerations suggest $r \geq 0.01$ (Boyle et al., 2006). Thus we, along with many other cosmologists, are interested in detecting tensor perturbations from inflation and measuring r from CMB observations.

1.2 The Cosmic Microwave Background

The CMB is relic radiation from when the Universe was younger and hotter. The CMB was created when the formation of Hydrogen made the Universe transparent to photons. By observing this radiation, cosmologists learn about the conditions of the early Universe. The small anisotropy in the radiation intensity has constrained cosmological models. The CMB polarization is the current frontier for probes of inflation because it has the smallest

contaminating signals.

The CMB was formed at redshift $z \approx 1100$ when the temperature became low enough (3000 K) for the ionization fraction to decrease rapidly⁹. Once the Universe was un-ionized, the photon mean free path became much greater than the Hubble volume. Most CMB photons last scattered at the time of recombination¹⁰. Thus the measured properties of the CMB directly correspond to early Universe physics. The CMB has a blackbody spectrum (Fixsen et al., 1996) with temperature $T_{\text{cmb}} = 2.72548 \pm 0.00057$ K (Fixsen, 2009). Spatially, the CMB is nearly uniform with $\Delta T/T \approx 10^{-5}$ (Smoot et al., 1992). These small anisotropies are the result of the initial perturbations from inflation modulated by plasma physics at recombination (Hu & Dodelson, 2002; Scott & Smoot, 2006).

1.2.1 CMB Measurements and Their Implications

CMB measurements, particularly of the temperature anisotropy, are one of the observational foundations of the standard cosmological model. The basic result of a temperature anisotropy experiment is a map of the CMB temperature $T(\theta, \phi)$. To extract the cosmologically relevant information, cosmologists typically decompose the map in spherical harmonics

$$\sum_{\ell,m} a_{\ell m} Y_{\ell m}(\theta, \phi) = T(\theta, \phi). \quad (1.24)$$

If the fluctuations are Gaussian and isotropic then all cosmological information is contained in the power spectrum

$$C_{\ell} \equiv \langle |a_{\ell m}^2| \rangle. \quad (1.25)$$

At low ℓ ($\lesssim 100$) the Sachs–Wolfe effect (Sachs & Wolfe, 1967) dominates, and $C_{\ell} \ell(\ell + 1)$ is nearly constant. At $\ell \gtrsim 100$, the oscillations of the photon–baryon fluid create acoustic

9. Eqs. 9.9, 9.29, and 9.37 of Ryden (2002)

10. Reionization at $z \approx 10$ caused 10% (Komatsu et al., 2011) of photons to scatter again.

peaks. At $\ell \gtrsim 1000$, the non-zero width of the last-scattering surface damps the anisotropies.

Cosmologists can only observe a single realization of the sky. At each ℓ there are only $2\ell + 1$ modes. Thus any measurement of the ensemble average C_ℓ has an irreducible sample variance¹¹ $\frac{2}{2\ell+1}C_\ell^2$. Moreover, practical¹² experiments only cover a fraction of the sky, $f_{\text{sky}} < 1$. Limited sky fraction reduces the effective number of modes measured so the sample variance increases to $\frac{1}{f_{\text{sky}}}\frac{2}{2\ell+1}C_\ell^2$ (Scott et al., 1994). Sample variance fundamentally limits the cosmological information that can be extracted from the CMB.

Existing CMB observations (Brown et al., 2009; Chiang et al., 2010; Larson et al., 2011; Reichardt et al., 2011; Keisler et al., 2011; Dunkley et al., 2011) combined with other observations (especially, galaxy clusters, weak gravitational lensing, supernova luminosity distances, and baryon acoustic oscillations (BAO). See e.g. Komatsu et al. (2011) for combined analyses) have shown that

1. The Universe is nearly flat. Komatsu et al. (2011) measured the curvature fraction $-0.0133 < \Omega_k < 0.0084$.
2. The initial perturbations were Gaussian and adiabatic (within measurement uncertainty).
3. The spectral index n_s is almost, but not exactly 1.
4. The initial perturbations at a given wavelength were in phase (Dodelson, 2003).
5. The Universe became un-ionized at $z \approx 1100$ and reionized at $z \approx 10$.
6. The current energy of the Universe is in dark energy ($72.8^{+1.5\%}_{-1.6\%}$) and dark matter ($22.7 \pm 1.4\%$).

11. Also called “cosmic variance.”

12. Ground-based experiments rarely cover the full sky. Even space-based experimenters reduce the effective fraction of the sky to avoid strong Galactic emission.

7. Structure formed from small initial perturbations which grew because of gravitational collapse.

These observations form the basis for the standard model of cosmology. Particularly, inflation predicts 1–4. To provide even more compelling evidence, we searched for the signal inflation generates in the CMB polarization.

1.2.2 CMB Polarization Measurements

The CMB is partially polarized (Hu & White, 1997; Hu & Dodelson, 2002). During decoupling, the quadrupole anisotropy in the local temperature led to a net linear polarization of the scattered radiation (Figure 1.1). In Thomson scattering, incoming radiation polarized parallel to the outgoing direction does not contribute to scattering. If the incoming radiation were isotropic, there would be no net polarization. Because the incoming radiation had a quadrupole anisotropy, the intensity was different in the two outgoing polarization states.

The Stokes parameters

$$I = LL^* + RR^* \quad (1.26)$$

$$Q = 2\Re[L^*R] \quad (1.27)$$

$$U = -2\Im[L^*R] \quad (1.28)$$

$$V = LL^* - RR^* \quad (1.29)$$

quantify the polarization properties of an electric field with components E_x and E_y . For later convenience I have introduced circular polarization components

$$L = \frac{E_x + iE_y}{\sqrt{2}} \quad (1.30)$$

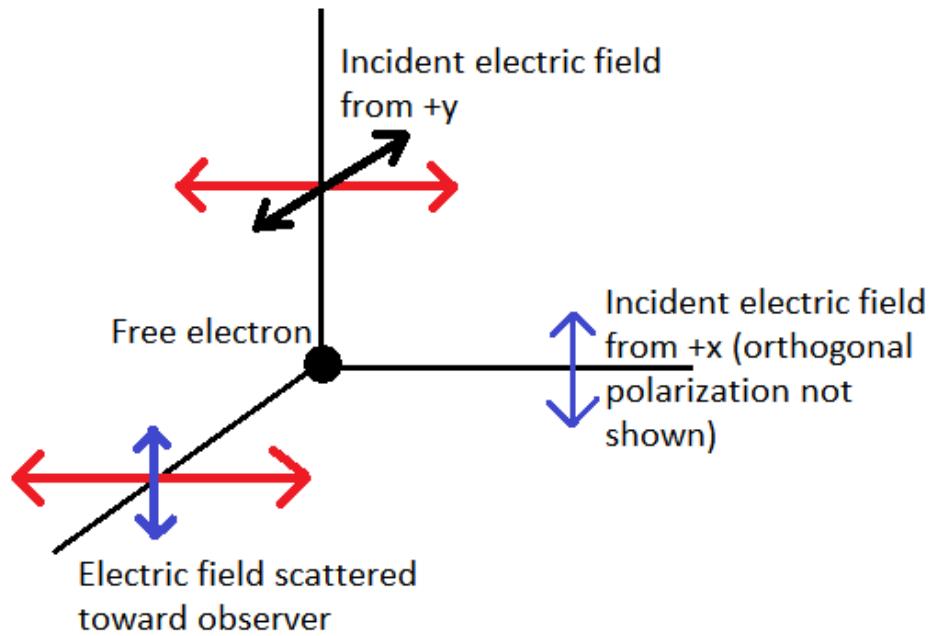


Figure 1.1: Quadrupole temperature anisotropies at decoupling created linear polarization in the CMB. Because electromagnetic waves are transverse, only one polarization state from the y direction can scatter toward us (z direction). Therefore if there was a difference in the electric field amplitudes in the x and y directions (equivalent to a radiation temperature difference), the scattered radiation will be polarized even though the incident radiation was not.

and

$$R = \frac{E_x - iE_y}{\sqrt{2}}; \quad (1.31)$$

although the net circular polarization (V) is negligible for the CMB in theory (Cooray et al., 2003). The Stokes parameters depend on the choice of coordinates¹³. If the coordinate system is rotated by an angle ψ so that $x \rightarrow x'$ and $y \rightarrow y'$,

$$Q' = Q \cos 2\psi + U \sin 2\psi \quad (1.32)$$

$$U' = -Q \sin 2\psi + U \cos 2\psi. \quad (1.33)$$

The quantity

$$(Q \pm iU)' = e^{\mp 2i\psi} (Q \pm iU) \quad (1.34)$$

has spin ± 2 , and I decomposed it into

$$Q \pm iU = \sum_{\ell,m} a_{\ell m}^{\pm} Y_{\ell m}^{\pm} \quad (1.35)$$

with spin-2 spherical harmonics Y^{\pm} . (The superscripts in a^{\pm} are part of the names of the expansion coefficients, not exponents.) The linear combinations

$$a_{\ell m}^E = -\frac{a_{\ell m}^+ + a_{\ell m}^-}{2} \quad (1.36)$$

$$a_{\ell m}^B = i\frac{a_{\ell m}^+ - a_{\ell m}^-}{2} \quad (1.37)$$

have definite parity: a^E (E mode) has even parity and a^B (B mode) has odd parity. The

13. This discussion follows Zaldarriaga & Seljak (1997).

corresponding power spectra are

$$C_\ell^{EE} = \langle a_{\ell m}^{E*} a_{\ell m}^E \rangle \quad (1.38)$$

$$C_\ell^{BB} = \langle a_{\ell m}^{B*} a_{\ell m}^B \rangle \quad (1.39)$$

$$C_\ell^{EB} = \langle a_{\ell m}^{E*} a_{\ell m}^B \rangle. \quad (1.40)$$

Since E and B modes have different parity, $C_\ell^{EB} = 0$ unless there is parity violation. For the same reason, primordial scalar perturbations do not produce B modes. However, tensor perturbations produce both E and B modes. Since r is small, scalar perturbations gave the dominant contribution to C_ℓ^{EE} . Therefore C_ℓ^{BB} is the most promising for a search for inflationary tensor perturbations.

Many experiments have measured non-zero C_ℓ^{EE} (Kovac et al., 2002; Sievers et al., 2007; Bischoff et al., 2008; Brown et al., 2009; Chiang et al., 2010; QUIET Collaboration et al., 2011), and the results agreed with the standard cosmological model. No experiment has detected B modes. The best upper limit on inflationary B modes is $r < 0.72$ (Chiang et al., 2010). Our main science goal in QUIET was to improve this limit. Astrophysical foregrounds¹⁴ make it extremely difficult to detect $r \lesssim 0.001$ (Dunkley et al., 2009; Smith et al., 2009).

1.3 The QUIET Experiment

QUIET (Buder, 2010d; QUIET Collaboration, 2012b) measured the CMB polarization to limit the B-mode signal from inflation. The QUIET instrument was a correlation polarimeter designed for high sensor density and low systematic errors. We completed two seasons of observations with this instrument; this thesis describes results from the second season.

14. The most important are synchrotron and dust emission from the Galaxy and gravitational lensing of the CMB.

The QUIET instrument (Figure 1.2) was a correlation polarimeter based on coherent-amplification technology. A 1.4-m side-fed Dragonian telescope collected the CMB radiation. An array of feed horns at the telescope focus formed nearly Gaussian beams for each of 90 pixels¹⁵. At the end of each feed horn, the radiation entered a septum polarizer that separated it into two circular polarization components (L and R). We used high-electron-mobility transistor (HEMT) based polarimeter modules to convert the two circular polarization components into a measurement of the Stokes linear polarization parameters Q and U . We cooled the feed horns, septum polarizers, and detectors to 20 K. We built a custom electronics system to bias the module active components and record the data. (“Receiver” means the cryostat, the components inside it, and the associated electronics hardware.)

We observed with this instrument from the Chajnantor Plateau, Chile from October 2008 until December 2010. In the first season (October 2008–June 2009) we measured the radiation in the Q band (centered at 43 GHz). [Bischoff \(2010\)](#); [QUIET Collaboration et al. \(2011\)](#); [Brizius \(2011\)](#); [Chinone \(2011a\)](#) described this measurement including the resulting limit $r < 2.2$. In the second season (August 2009–December 2010) we measured the radiation in the W band (centered at 94.5 ± 0.8 GHz with 10.7 ± 1.1 GHz bandwidth). Each receiver (Q or W) worked at only one frequency band, and we could only observe with one receiver at a time. [QUIET Collaboration \(2012a\)](#) summarizes the CMB polarization results from the second season; this thesis describes them in more detail. Details about the instrument, observations, analysis, and results apply to W band unless otherwise noted.

This thesis contains the following elements. Chapter 1 introduces QUIET. Chapter 2 describes the electronics system. Chapter 3 describes data-management software and practices. Chapter 4 describes the W-band observation from Chile. Chapter 5 describes the data

15. “Module” means the polarization detectors associated with each pixel and, in a more general sense, the corresponding feed horn, septum polarizer, data, etc. We attached six modules to orthomode transducers that made them sensitive to temperature anisotropies; I excluded these modules (“differential-temperature”) from analysis in this thesis. Six of the remaining 84 modules were not functional or had no data passing cuts (§5.3.2). An additional four modules had a non-functional diode. Thus the total diode yield for polarization was $308/336 = 92\%$.



Figure 1.2: The QUIET instrument deployed to the Chajnantor Plateau in Chile's Atacama Desert to reduce atmospheric emission.

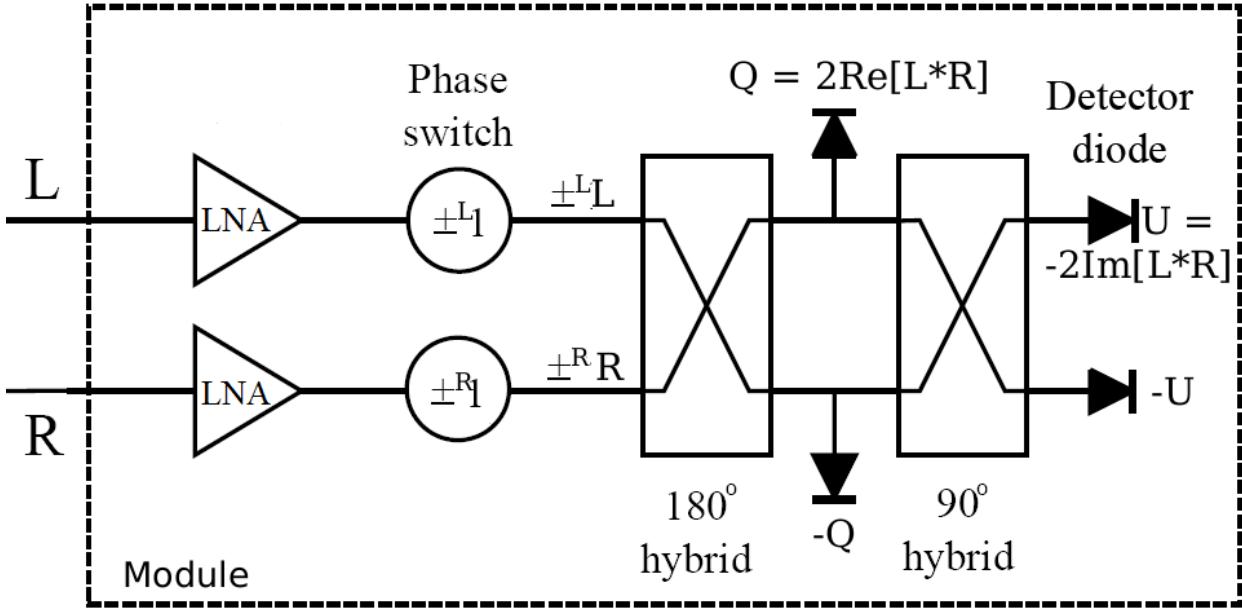


Figure 1.3: The QUIET module converted L and R into measurements of Q and U . We modulated each measurement twice: once with each phase switch.

analysis. Chapter 6 describes the results and my conclusions.

1.3.1 Module Operation Principles

The QUIET module implemented our polarization-detection scheme. Here I outline the module operation; §1.3.3 gives implementation details. The following discussion applies to an “ideal” module—one with no systematic errors. Bischoff (2010); QUIET Collaboration (2012b); and §1.3.3 of this thesis describe module non-idealities. The module was a pseudo-correlation polarimeter (Gaier, 2003): the polarization signal appeared as the product of two independently-processed polarization states (Figure 1.3). The module operated entirely at radio frequency (RF). There was no intermediate frequency or local oscillator. The final stage of the module rectified and integrated the signal.

The module had two input ports. Normally, these received the left and right circular-polarization components (L and R) of the CMB. Each input fed into a “leg” of the module. Each leg contained a phase switch that applied 0° or 180° phase shift to the incoming

radiation. I denote the phase switch outputs by $\pm^L L$ and $\pm^R R$ ¹⁶. The two legs then merged in a 180° hybrid coupler. The coupler output $\frac{\pm^L L \pm^R R}{\sqrt{2}}$ and $\frac{\pm^L L \mp^R R}{\sqrt{2}}$.

Detector diodes¹⁷ (Q1 and Q2) rectified half of each signal with resulting voltages

$$\begin{aligned} V_{Q1} &\propto \frac{\pm^L L \pm^R R \pm^L L^* \pm^R R^*}{2} = \frac{LL^* \pm^L \pm^R (L^* R + LR^*) + RR^*}{4} \\ &= I/4 \pm^L \pm^R \Re[L^* R]/2 = \frac{I \pm^L \pm^R Q}{4} \end{aligned} \quad (1.41)$$

$$\begin{aligned} V_{Q2} &\propto \frac{\pm^L L \mp^R R \pm^L L^* \mp^R R^*}{2} = \frac{LL^* \pm^L \mp^R (L^* R + LR^*) + RR^*}{4} \\ &= \frac{I - \pm^L \pm^R Q}{4} \end{aligned} \quad (1.42)$$

after using Eqs. 1.26 and 1.27. The other halves of the signal entered a second, 90° hybrid coupler which output $\frac{\pm^L L \mp^R R - i(\pm^L L \pm^R R)}{\sqrt{8}}$ and $\frac{\pm^L L \mp^R R + i(\pm^L L \pm^R R)}{\sqrt{8}}$. Two more detector diodes (U1 and U2) rectified these outputs with resulting voltages

$$\begin{aligned} V_{U1} &\propto \frac{\pm^L L \mp^R R - i(\pm^L L \pm^R R)}{\sqrt{8}} \frac{\pm^L L^* \mp^R R^* + i(\pm^L L^* \pm^R R^*)}{\sqrt{8}} \\ &= \frac{1}{8} \left(LL^* - \pm^L \pm^R LR^* + iLL^* + \pm^L \pm^R iLR^* \right. \\ &\quad \left. - \pm^L \pm^R RL^* + RR^* - \pm^L \pm^R iRL^* - iRR^* \right. \\ &\quad \left. - iLL^* + \pm^L \pm^R iLR^* + LL^* + \pm^L \pm^R LR^* \right. \\ &\quad \left. - \pm^L \pm^R iRL^* + iRR^* + \pm^L \pm^R RL^* + RR^* \right) \\ &= \frac{1}{8} \left(2LL^* + \pm^L \pm^R 2iLR^* + 2RR^* - \pm^L \pm^R 2iRL^* \right) \\ &= \frac{I - \pm^L \pm^R 2\Re[iL^* R]}{4} = \frac{I - \pm^L \pm^R 2\Im[L^* R]}{4} \\ &= \frac{I - \pm^L \pm^R U}{4} \end{aligned} \quad (1.43)$$

16. The two phase switches can operate independently so there are four possible choices for the \pm signs, denoted by “ \pm^L ” and “ \pm^R .”

17. I use “diode” to refer to any of the detector diodes or their associated data.

$$\begin{aligned}
V_{U2} &\propto \frac{\pm^L L \mp^R R + i(\pm^L L \pm^R R)}{\sqrt{8}} \frac{\pm^L L^* \mp^R R^* - i(\pm^L L^* \pm^R R^*)}{\sqrt{8}} \\
&= \frac{1}{8} \left(LL^* - \pm^L \pm^R LR^* - iLL^* - \pm^L \pm^R iLR^* \right. \\
&\quad - \pm^L \pm^R RL^* + RR^* + \pm^L \pm^R iRL^* + iRR^* \\
&\quad + iLL^* - \pm^L \pm^R iLR^* + LL^* + \pm^L \pm^R LR^* \\
&\quad \left. + \pm^L \pm^R iRL^* - iRR^* + \pm^L \pm^R RL^* + RR^* \right) \\
&= \frac{1}{8} \left(2LL^* - \pm^L \pm^R 2iLR^* + 2RR^* + \pm^L \pm^R 2iRL^* \right) \\
&= \frac{I \pm^L \pm^R 2\Re[iL^*R]}{4} = \frac{I - \pm^L \pm^R 2\Im[L^*R]}{4} \\
&= \frac{I \pm^L \pm^R U}{4} \tag{1.44}
\end{aligned}$$

after using Eqs. 1.26 and 1.28.

The electronics (§2) modulated the L phase switch at 4 kHz and the R phase switch at 50 Hz. After digitization we recorded the average (“TP”) and demodulation (“DE”) of the data in both 4-kHz phase switch states. For the Q1 diode these are

$$TP \propto (I \pm^L \pm^R Q) + (I \mp^L \pm^R Q) \propto I \tag{1.45}$$

$$DE \propto (I \pm^L \pm^R Q) - (I \mp^L \pm^R Q) \propto \pm^R Q. \tag{1.46}$$

We stored the data at 100 Hz so adjacent data samples had opposite signs of \pm^R . By subsequently differencing two adjacent samples we obtained “double-demodulated” (DD) data $\propto Q$. (The Q2 diode yielded the same except with $Q \rightarrow -Q$. The U diodes yielded the same except $Q \rightarrow U$.) Double demodulation reduced the effect of amplifier gain variations (1/f noise) and eliminated instrumental polarization otherwise caused by gain mismatch between the two legs (QUIET Collaboration, 2012b). Furthermore each module (and therefore each

pixel) fully characterized the linear polarization.

1.3.2 Optics

Our optical system consisted of the telescope, ground screen, feed horns, and septum polarizers. We chose a 1.4-m side-fed Dragonian antenna for the telescope. A comoving, absorbing ground screen limited the sidelobes. We used corrugated feed horns to form the beams for each module. Each feed horn sent its radiation to a septum polarizer, which separated L and R into different waveguides for the module.

The telescope used a two-reflector design ([Imbriale et al., 2011](#)) that satisfied the Mizuguchi condition ([Mizugutch et al., 1976](#)). The result had a wide field of view (Figure 1.4), low instrumental polarization, and low sidelobes. We made the two mirrors out of aluminum, the backs of which we light-weighted and connected by hexapods with adjustable turnbuckles to a steel support structure. This support structure attached the telescope (and the remainder of the instrument) to the existing Cosmic Background Imager (CBI) mount ([Padin et al., 2002](#)). The mount had three rotation axes (azimuth, elevation, and “deck”). The azimuth axis was capable of fast scanning ($\approx 5^\circ/\text{s}$) so that the typical scan frequency was well above the typical $1/\text{f}$ knee frequency. By regularly (§4.2.1) rotating the third axis (deck), we modulated the instrumental-polarization axis, thereby suppressing instrumental polarization.

We enclosed the optics with a comoving, ambient-temperature, absorbing ground screen. The ground screen had two main parts: lower (LGS) and upper (UGS). The LGS was an aluminum box surrounding the mirrors with a hole for the cryostat. The UGS was a telescoping cylinder that surrounded the path of the main beam. Several smaller pieces formed the bottom ground screen (BGS) that covered the floor of the LGS and the gap between the LGS and the cryostat. We coated the entire inner surface of the ground screen with microwave absorber Emmerson Cummings HR-10 Eccosorb ([Brizius, 2011](#)). We sealed the Eccosorb with Volara (a microwave-transparent polyethylene foam) for weatherproofing.

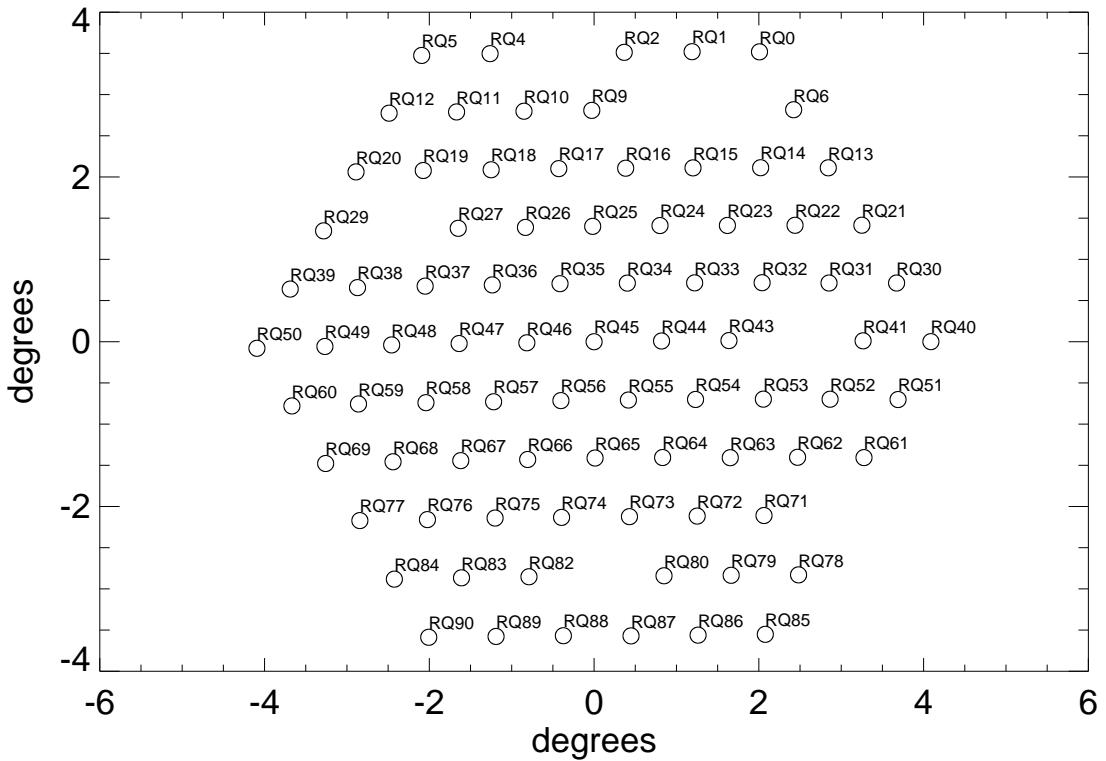


Figure 1.4: The module array had a large (8°) field of view on the sky. Each circle represents one module beam with diameter equal to the FWHM of $11.66'$. The modules and their corresponding positions are labeled RQ0–RQ90. Modules RQ7, 8, 28, and 42 did not function. Modules RQ3 and RQ70 did not have any data passing selection. There was no module at position RQ81.

We coated the outer surface with white paint to minimize radiative loading.

Due to manufacturing delays, we did not install the UGS until January 2010. As a result, two significant sidelobes were present in the 2009 data. The “triple-reflection” sidelobe was due to a portion of the beam making an additional, unintended reflection off the secondary mirror. The result was a narrow (1°) sidelobe 50° away from the main beam (Chinone, 2011b). The “spillover” sidelobe resulted from part of the main beam escaping above the secondary mirror and LGS (Figure 1.5). The spillover appeared at $\approx 50^\circ$ from the main beam in the opposite direction from the triple-reflection sidelobe; however, the spillover sidelobe was much more extended ($\pm 20^\circ$ in the direction towards the main beam). Holes in the BGS caused additional spillover (Sanders, 2010) (see also §4.4). We mitigated these effects with data selection (§5.3.2).

A corrugated, linear-flare feed horn formed the beam of each module. We constructed these feed horns as a hexagonally-close-packed “platelet array” (Gundersen & Wollack, 2009). We machined many large aluminum plates so that each plate formed a slice of all the feed horns, transverse to the optical axis. We then stacked the plates along the optical axis and diffusion bonded them together to form a single component containing all the feed horns. The platelet design incorporated light-weighting holes that also functioned as access holes for attachment screws. The fabrication costs of these arrays were at least an order of magnitude less than for the same number of electroformed feed horns. Measurements of the return loss and beam characteristics of the platelet arrays showed performance comparable to an equivalent electroformed design. The feed-horn full-width half-maximum (FWHM) was $7.0\text{--}8.6$ across the array, and instrumental polarization was -29 dB or better.

Radiation exited each feed horn into a septum polarizer (Bornemann & Labay, 1995; Bannister et al., 2006), which separated L and R for input to the QUIET detector modules. The polarizer itself was a square waveguide with a stepped-thickness septum in the center.

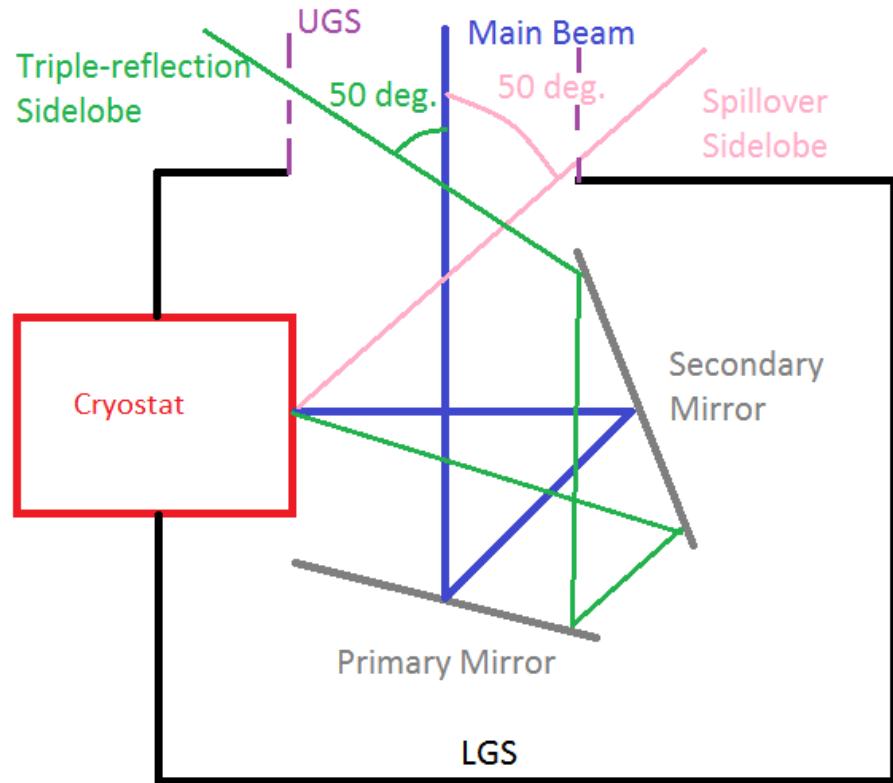


Figure 1.5: Before installation of the upper ground screen (in the middle of the observing season) two significant far sidelobes escaped the ground screen. The upper ground screen captured both sidelobes on a microwave absorber.

The presence of the septum added a 90° phase shift to the TE_{01} mode in the waveguide¹⁸. A waveguide splitter followed the polarizer and physically separated the two outputs into rectangular WR-10 waveguides with spacing compatible with the module waveguide inputs. Although these waveguides carried components corresponding to circular polarization on the sky, the radiation in each waveguide was in a single, linearly-polarized (TE) mode. Polarizer imperfections caused temperature to polarization ($I \rightarrow Q/U$) leakage¹⁹. Our lab measurements (Nixon, 2008) implied leakages of $\approx 0.3\%$.

1.3.3 Modules

The QUIET detectors were miniaturized polarimeter modules (Cleary, 2010) that converted L and R into Q and U as described in §1.3.1. Each module was based on HEMT amplifiers. Previous experiments including PIQUE (Hedman et al., 2001), DASI (Leitch et al., 2002), CBI (Padin et al., 2002), *WMAP* (Jarosik et al., 2003), and CAPMAP (Barkats et al., 2005) also used HEMT-based polarimeters, but we used advances in millimeter-wave circuit technology and packaging (Gaier et al., 2003) to replace waveguide-block components and connections with strip-line-coupled devices. The resulting modules were $3.2\text{ cm} \times 2.9\text{ cm}$ (Figure 1.6), almost an order of magnitude smaller than a comparable waveguide design. Each module had two waveguide inputs matched to the two septum-polarizer outputs. A probe in each port coupled the radiation into ($50\text{-}\Omega$) microstrip transmission lines. The first active component in each leg was a HEMT-based monolithic-microwave-integrated-circuit (MMIC) amplifier²⁰. A second MMIC provided additional amplification before the phase switch. After the phase switch, the signal passed through a bandpass filter before the third (final) MMIC. After the final MMIC, the two legs combined in the hybrid coupler. (For

18. §2.2.3 of Brizius (2011)

19. Appendix 10.3 of QUIET Collaboration (2012b)

20. Herein, “MMIC” appearing alone means one of the QUIET MMIC amplifiers.

conceptual simplicity, §1.3.1 separates the hybrid coupler into a first 180° coupler and second 90° coupler. We manufactured both couplers as a single six-port device.) Each coupler output (Q1, Q2, U1, and U2) passed through a second bandpass filter before rectification at its detector diode.

The MMICs were low-noise amplifiers with four HEMT stages manufactured in an Indium-Phosphide (InP) process (QUIET Collaboration, 2005). Each MMIC provided ≈ 25 dB gain with noise temperature 50–80 K when cooled to a physical temperature of 20 K²¹. At 20 K the MMICs required ≈ 10 mA of drain current bias, 0.5 V drain voltage, and 0.1 V gate voltage for optimum performance. Because the number of module electrical connections was limited, we connected MMIC drains together²² so that each leg had three gate connections and two drain connections. Thus each module had 10 MMIC bias connections in total. We optimized the sensitivity over these 10 parameters with a sparse-wire-grid calibration source (Tajima et al., 2012) and downhill-simplex algorithm²³.

The phase switches were two-path circuits also fabricated in InP. One of the paths had an additional half-wavelength length so that radiation passing through it acquired a 180° phase shift. A PiN (p-type, intrinsic, n-type) diode controlled whether radiation could pass through each path. Reverse biasing the diode put it into a high-impedance state, blocking that path. When forward biased to $400\ \mu\text{A}$, the transmission was \approx unity (Figure 2.6). Typically we forward biased one diode and reverse biased the other so the phase shift was either 0° or 180° as described in §1.3.1. However, we periodically (§4.2) made a diagnostic “offset” measurement by reverse biasing both diodes to block all RF power. Since each module had two phase switches and each switch had two diodes, there were four phase

21. The improvement in noise temperature below a physical temperature of 40 K was percent-level (Tajima, 2009b). Including all the active and passive components, the typical module noise temperature was 100 K. See Appendix F for the total array sensitivity.

22. The second and third MMIC drains were common. Some modules had two independent gate controls for the first MMIC and a single common gate for the second and third MMICs. Other modules had one gate control per MMIC.

23. §2.4.4 of Brizius (2011) and Becker (2008)

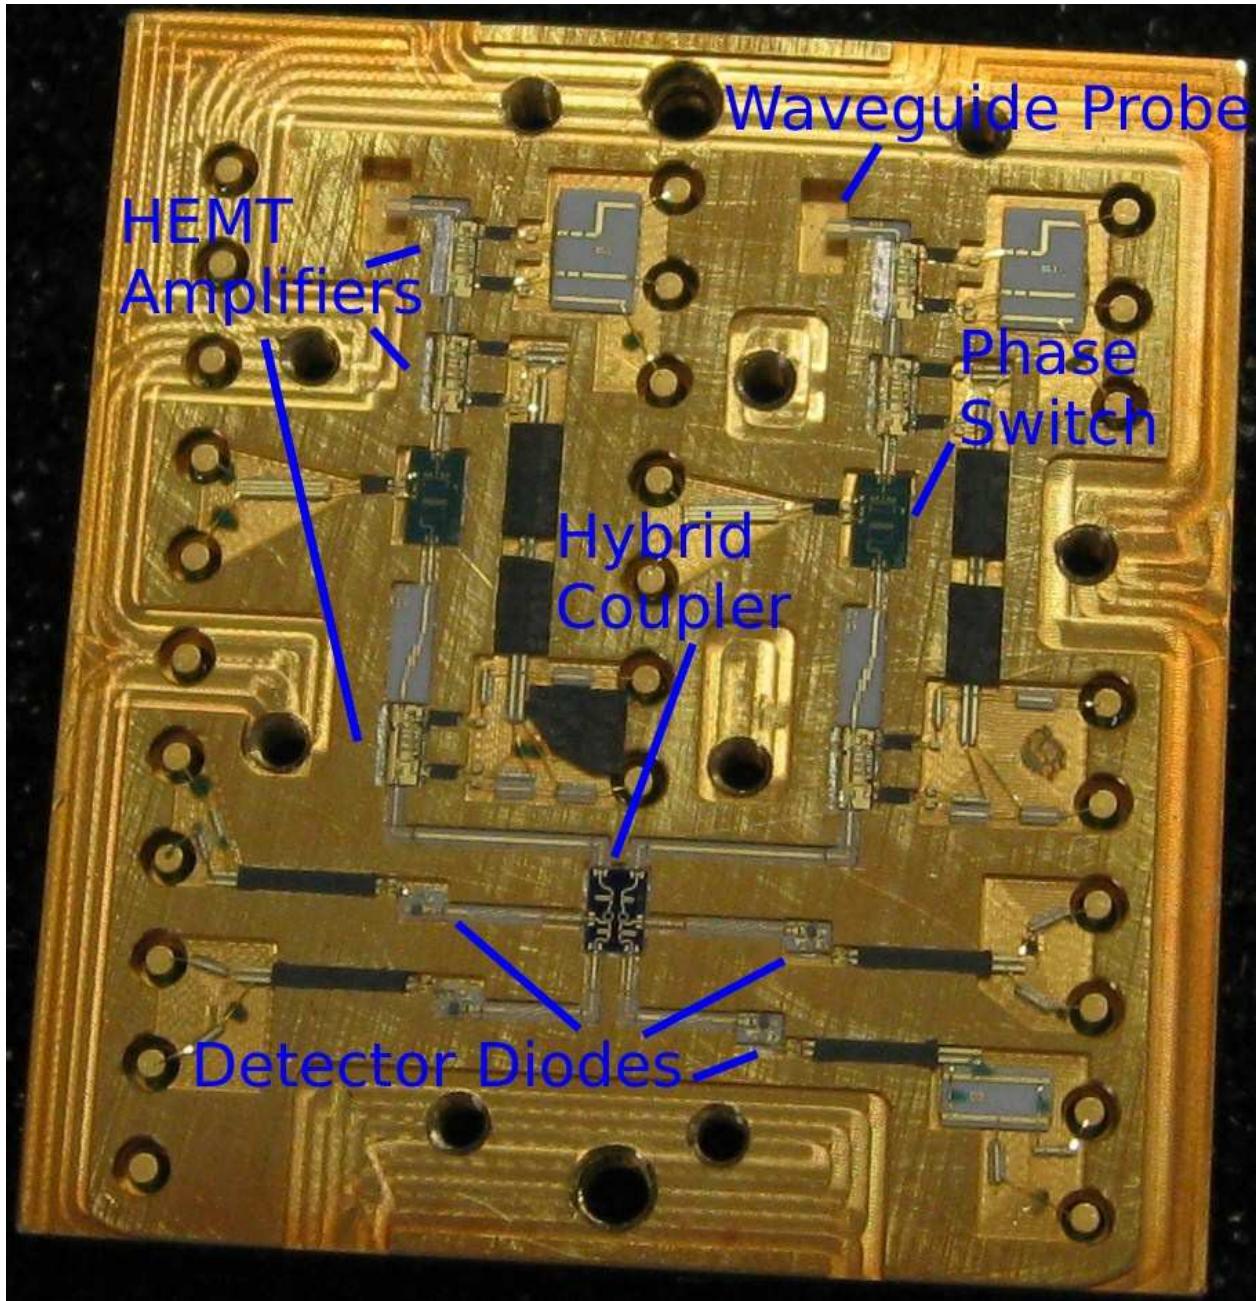


Figure 1.6: The QUIET module was a miniaturized correlation polarimeter. We used microstrip transmission lines to couple different components. This reduced the size by a factor of 10 compared to a traditional waveguide-based design.

switch bias currents per module. Differences in transmission between the two paths would cause $I \rightarrow Q/U$ leakage if we did not use double demodulation. Because we used double demodulation, we did not attempt to balance the transmission by adjusting the diode biasing.

The hybrid coupler was a passive, planar InP device that summed its two inputs with either 0° , 90° , or 180° phase shifts as described in §1.3.1. Our design used Lange couplers and Schiffman phase delay lines to perform these functions (QUIET Collaboration, 2005). Because the phase shifts were not exactly 90° or 180° , the detector angles were different from their ideal values²⁴ (Kusaka, 2007), and the noise correlations among the diodes also changed from their ideal values (Bischoff & Kusaka, 2009).

The detector diodes were Schottky diodes (Agilent HSCH-9161) which we operated in the square-law regime. At 20 K the diodes required bias ($\approx 30 \mu\text{A}$) to make the impedance reasonable (200Ω)²⁵. The diode bias and readout was differential so each diode required two connections for a total of eight diode connections per module.

A brass housing contained the components of each module. Inside the housing, microstrip transmission lines connected the RF components. Ribbon bonds connected DC bias and readout points to pins leading out of the module. Each module had 23 such pins: 10 MMIC, 4 phase switch, 8 diode, and 1 shared module ground. Since all external module connections were DC, the modules were easy to test and optimize in a repeatable way.

1.3.4 Cryostat

The cryostat mounted and supported the platelet array, septum polarizers, and modules (Figure 1.7). Radiation entered the cryostat through an ultra-high-molecular-weight—polyethylene window with an expanded-Teflon anti-reflection coating²⁶. With this coating

24. e.g. the Q diodes acquired some U sensitivity

25. See §3.1.7 of Bischoff (2010); Kapner (2006); Bischoff (2006b) for more details on diode biasing.

26. §2.6.5 of Newburgh (2010)

the loss through the window was at the percent level (The loss increased the system noise temperature by 4 K.) for microwave frequencies. A polystyrene filter blocked infrared radiation, reducing thermal loading on the cold stage. We made an interface plate²⁷ in the cold stage between the platelet array and septum polarizers; the interface plate included waveguide transitions between the circular waveguide at the end of each feed horn and the square waveguide at the input of each septum polarizer. We cooled the cold stage containing the modules and optical components to ≈ 25 K²⁸ using two CTI Cryogenics Cryodyne 1020 Gifford-McMahon refrigerators. We attached a heater to each cold head, and a temperature controller (Cryo-con model 32B, hereafter “CPID”) regulated the cold-head temperatures using a feedback loop. We put silicon diode thermometers on the detectors, platelet array, refrigerator heads, and secondary 80-K stage (The CTI 1020 has two cold stages. The 80-K stage cooled a radiation shield.) to continuously monitor the cooling and temperature regulation performance. The refrigerators had a mechanical cycle frequency of 1.2 Hz. I²⁹ found a response at this frequency in the DD data (§5.3.2). Styccast-epoxy-sealed hermetic pass-throughs allowed the bias and readout connections (§2.1) to be brought out of the cryostat.

27. Newburgh et al. (2006) and §2.3.3 of Brizius (2011)

28. We had difficulty keeping the temperature stable. The standard deviation of the module temperature during the season was 1 K. See §4.4 and Figure 1.8 for more details.

29. “I” indicates particular contributions of the author.

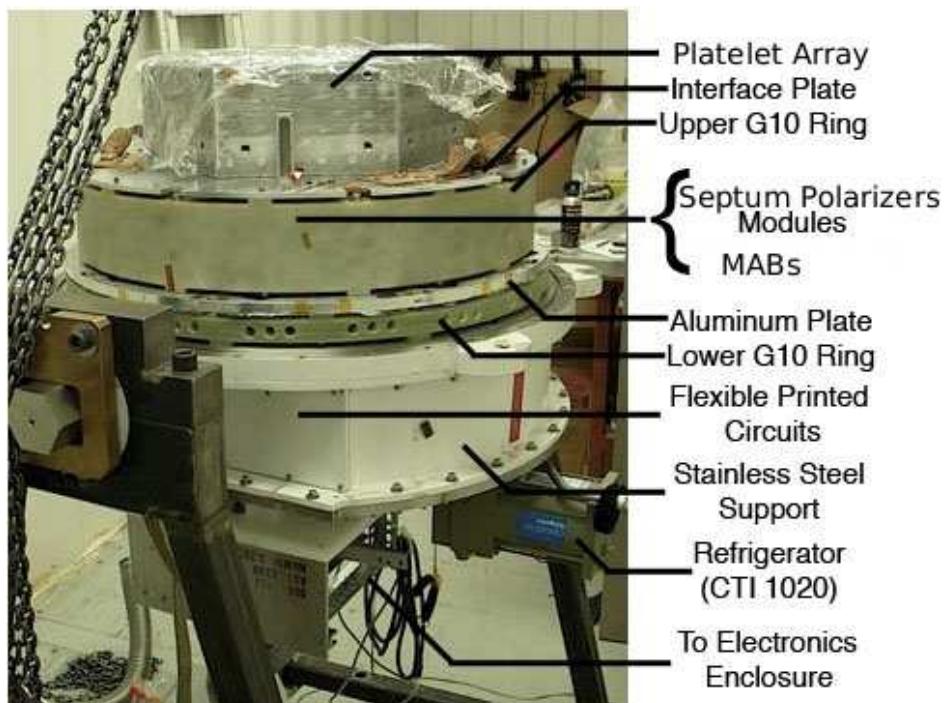


Figure 1.7: The cryostat contained the platelet array, septum polarizers, and modules. We cooled them to $\approx 25\text{ K}$ to reduce the system noise.

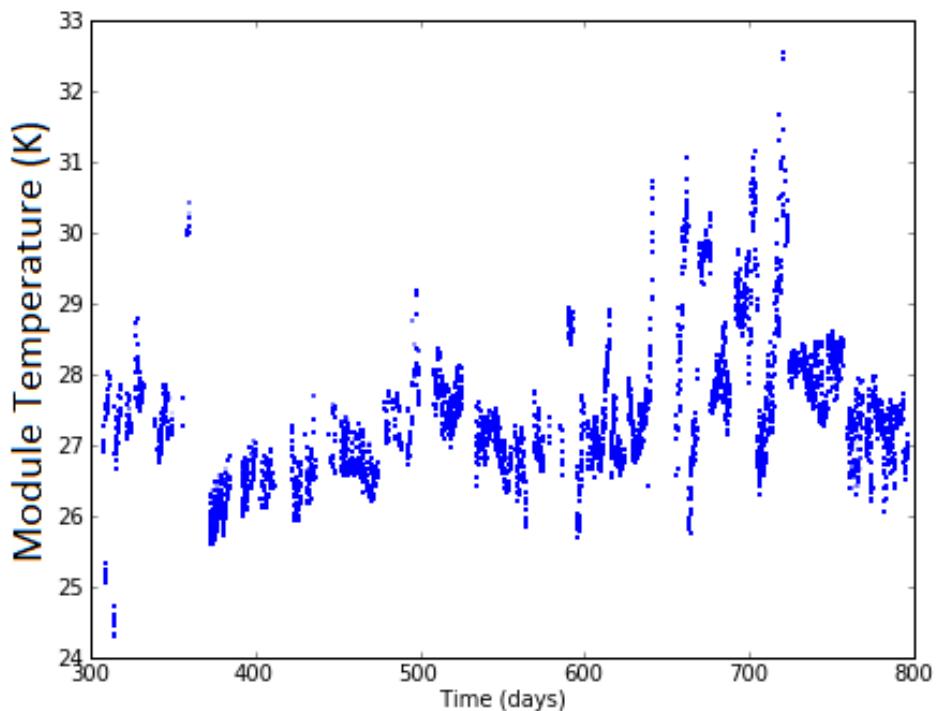


Figure 1.8: The cryostat temperature varied throughout the season due to problems with the cold heads and vacuum. The main effect was that the responsivity changed; we took it into account in the responsivity model and associated systematic error. (I show the temperature of the central module, RQ45.)

Chapter 2

Electronics

They said it couldn't be done, but sometimes it doesn't work out that way.

Casey Stengel

The electronics system provided module biasing, timing synchronization, and data acquisition (DAQ) for the receiver. We divided and assigned these tasks to four systems (Figure 2.1): (1) Passive Interfaces, (2) Bias, (3) Readout, and (4) Data Management. The Passive Interfaces system (§2.1) created an interface between the modules, the biasing electronics, and the DAQ. The Bias system (§2.2) generated the voltage or current necessary to operate each module active component. The Readout system (§2.3) amplified and digitized the module outputs. It included the timing synchronization hardware. The Data Management system (§3) commanded the other systems and recorded the data.

2.1 Passive Interfaces

The main function of the Passive Interfaces system was to bring the module electrical connections out of the cryostat. We accomplished this function in three stages: first, Module Assembly Boards (MAB) connected directly to the modules. Second, Flexible Printed Circuits (FPC) connected to the MABs and passed through the cryostat vacuum seal. Third, Array Interface Boards (AIB) interfaced between the FPCs and standard ribbon cables.

Each MAB was a printed circuit board with pin sockets for seven modules. (Sometimes “MAB” refers to the associated modules e.g. “MAB3 has high noise.”) We made five electrically identical MAB designs with different physical shapes as needed to fill the focal plane (Figure 2.2); there were 13 MABs in total. Each module had 27 pins (see §1.3.3 for

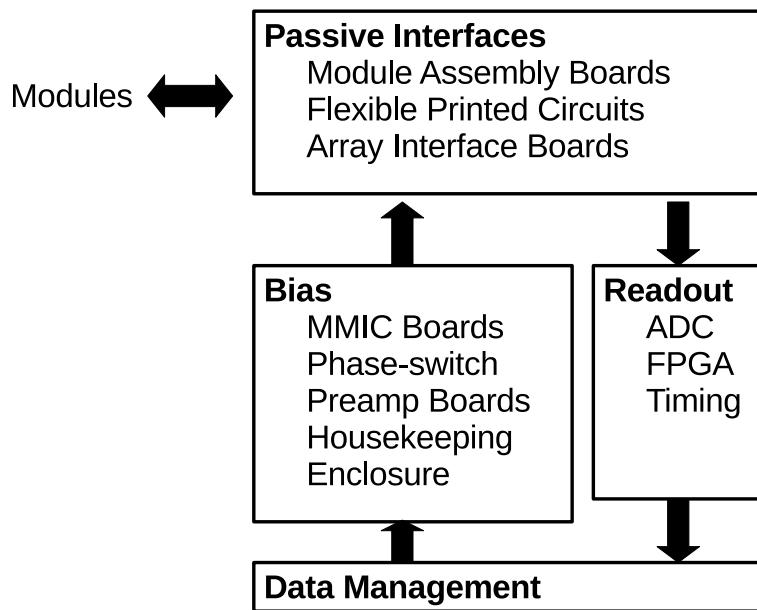


Figure 2.1: We divided the electronics into four systems: Passive Interfaces, Bias, Readout, and Data Management. Commands flowed from the observers, through the Bias and Passive Interfaces systems, and to the modules. The module outputs flowed through the Passive Interfaces system to the Readout system, which digitized them, and the Data Management system, which recorded them.

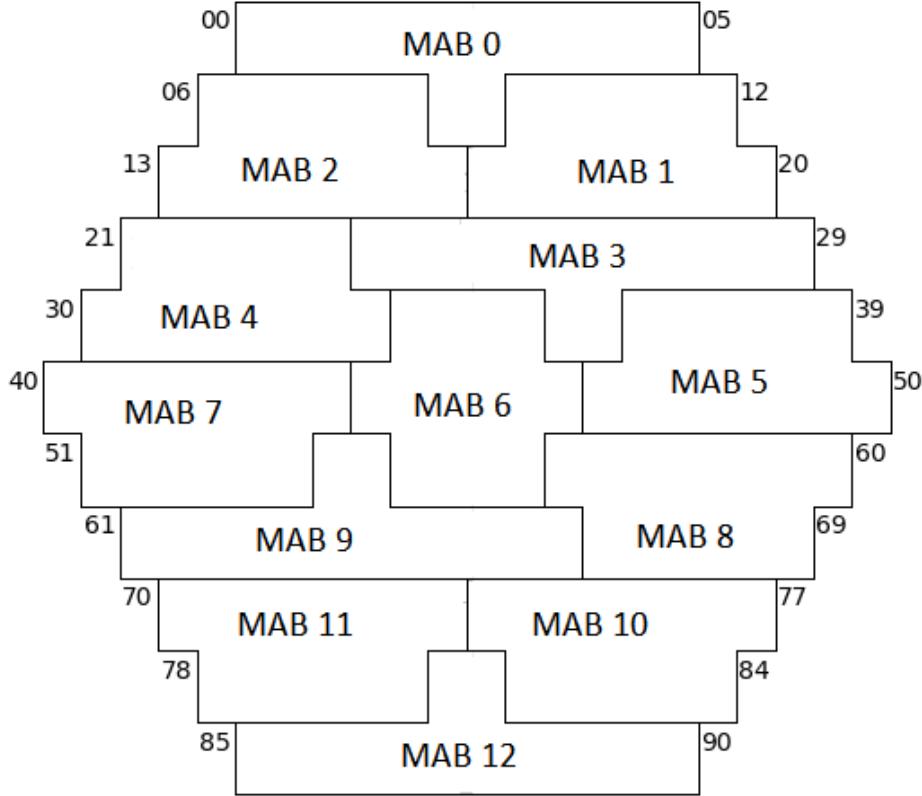


Figure 2.2: We tiled 13 MABs to fill the focal plane. Numbers to the left and right indicate module position numbers. MABs 0 and 3 had the same shape and were physically identical circuit boards. Similarly, MABs 1, 2, 4, and 5; 7, 8, 10, and 11; and 9 and 12 were physically identical. Thus we needed only five MAB circuit layouts. The modules could not rotate 180° in the focal plane so a duplicate MAB 4 cannot serve in place of MAB 7, for example.

their functions; four pins were unused). Voltage clamps and RC low-pass filters on the MAB protected the sensitive components inside the modules from damage caused by transients or accidental overvoltage¹. After protection, the MAB circuitry connected the (23 used) module pins to 40-contact Hirose connectors. Each MAB had 5 such connectors: 2 for MMIC bias pins, 2 for detector diodes, and 1 for phase switches. Each Hirose connector accepted one FPC.

1. See §2.5.1 of [Brizius \(2011\)](#); [Bischoff & Kapner \(2006\)](#); [Bischoff \(2007\)](#) for details about the MAB protection circuitry for each module component.

The FPCs were high-density (40 conductors each) 32"-long flexible circuit boards². One end of each FPC connected to an MAB. Each FPC carried either MMIC, diode, or phase-switch signals for that MAB³. We grouped the five FPCs for each MAB together for cable routing (there were $5 \times 13 = 65$ FPCs in total). We clamped each FPC group to the cryostat 80-K stage to reduce thermal loading. Each FPC group passed through the cryostat wall in a Stycast-epoxy-filled custom hermetic seal⁴. The other end of each FPC connected to an AIB.

The AIBs provided a second layer of electrical protection for the modules. We attached them to the back of the cryostat (Figure 2.3). When cold, the diodes comprising the voltage clamps on the MABs had drastically different I-V curves than when warm and were no longer effective at protecting the modules. (The filter resistances and capacitances could change as well.) Therefore, we duplicated the module protection circuitry on the AIBs. Each AIB accepted only MMIC, diode, or phase-switch FPCs. A MMIC AIB accepted four MMIC FPCs (from two MABs) and combined and routed the signals (after protection) to one AIB Cable (Harting) connector. The AIB Cables connected the AIBs to the bias boards (§2.2). We made each AIB Cable from two 80-pin shielded ribbon cables; AIB Cable Boards (Figure 2.4, see also [Dumoulin \(2007\)](#)) interfaced between the Harting connectors and the ribbon cables. A “Preamp” AIB (corresponding to “Preamp Boards,” §2.2) accepted four diode FPCs (from two MABs), and a Phase-switch AIB accepted three phase-switch FPCs (from three MABs). Each AIB connected one AIB Cable, which connected to one bias board (see Appendix D for the mapping between modules, MABs, and bias boards). We surrounded the AIBs and AIB Cable connections with a weatherproofing structure (Figure 2.5) to protect them from damage due to moisture and dust.

2. Figure 2.25 in §2.5.1 of [Brizius \(2011\)](#)

3. Table 3 of [Nixon & Dumoulin \(2006\)](#); [Nixon \(2006, 2007\)](#) list which signal is on which FPC conductor.

4. §4.3.2 of [Bischoff \(2010\)](#)

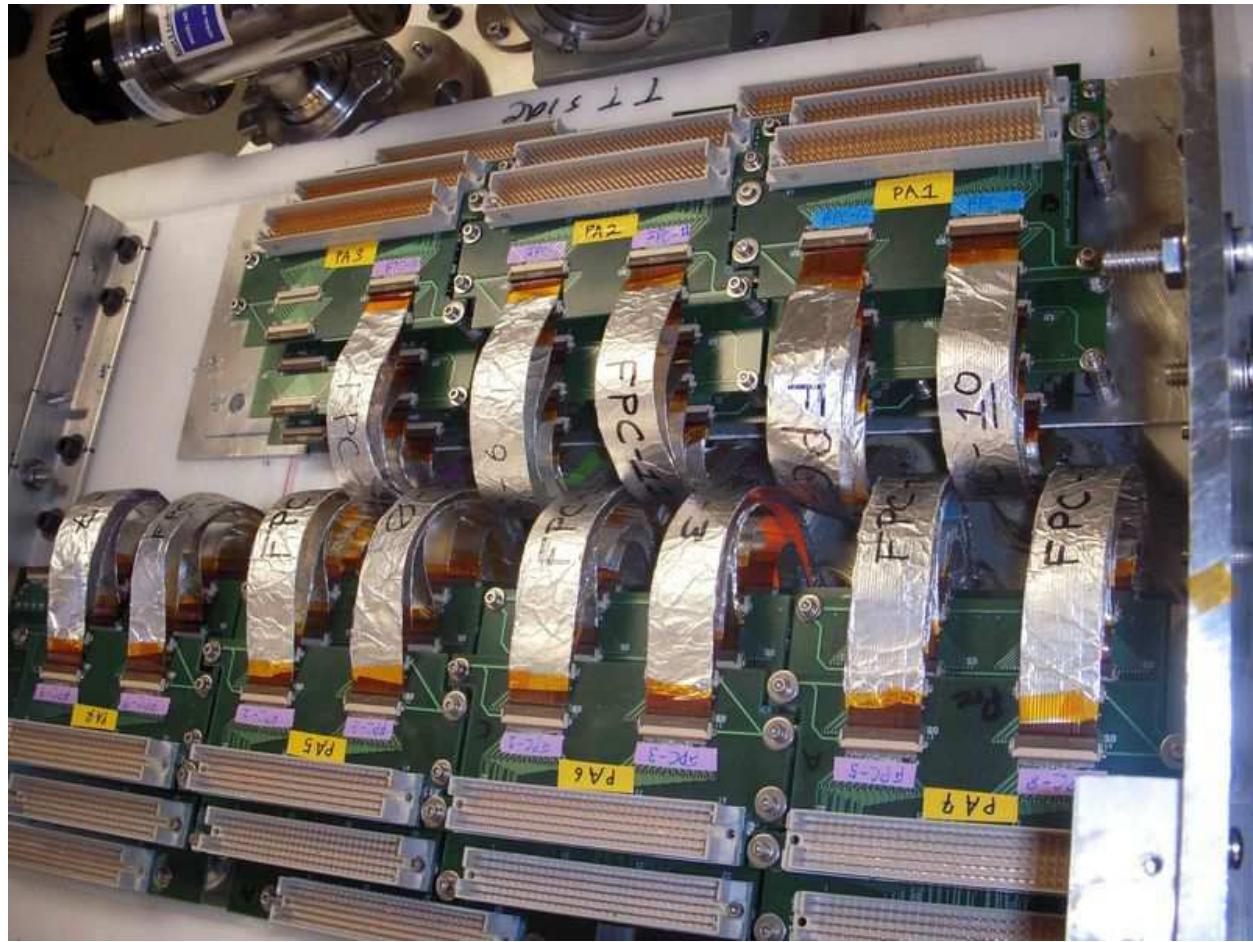


Figure 2.3: We attached the Array Interface Boards (AIB) to the back of the cryostat. Each AIB connected one bias board to two or three MABs. The AIBs also provided electrical protection for the modules. SOURCE: Alison Brizius

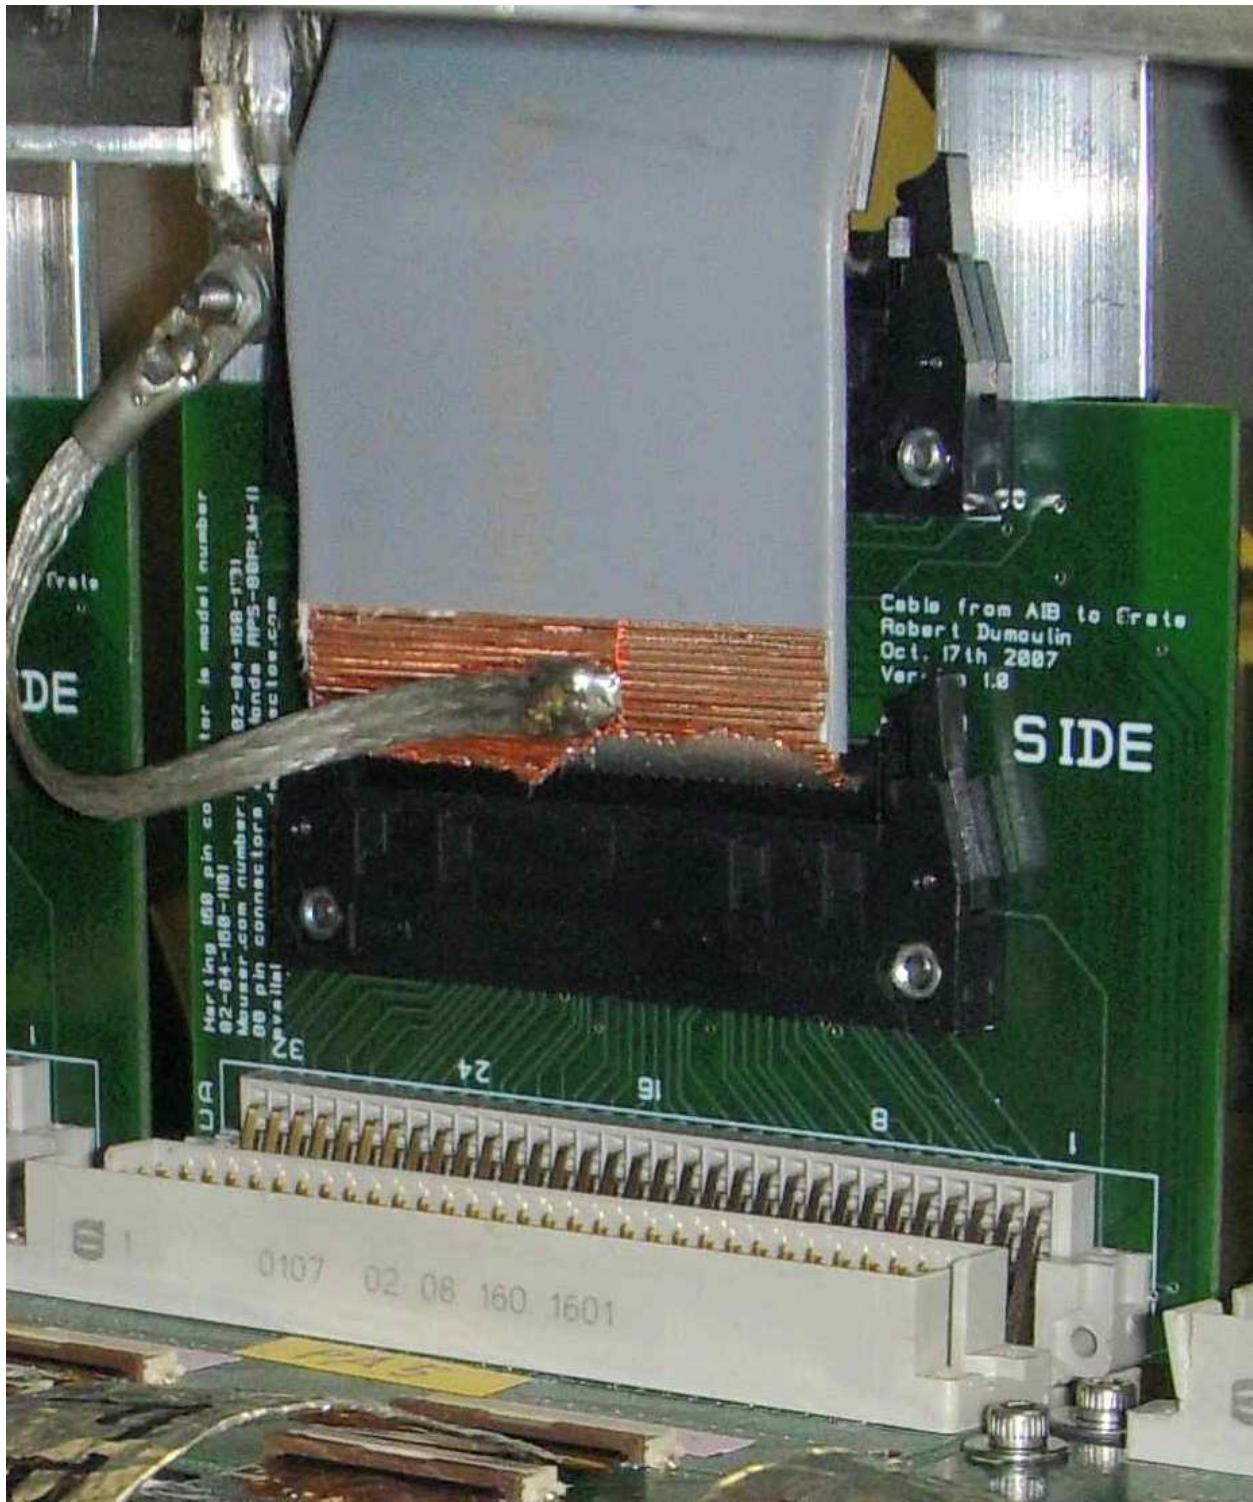


Figure 2.4: AIB Cable Boards interfaced between the Array Interface Board (AIB) Harting connectors (bottom) and the two ribbon cables making up each AIB Cable. SOURCE: Alison Brizius



Figure 2.5: Weatherproofing protected the Array Interface Boards (AIB) and AIB Cable connections.

2.2 Bias

We made custom integrated-circuit boards to provide the bias signals needed to operate the modules. We apportioned the module components that require biasing to three types of bias boards: MMIC, Phase-switch, and Preamp. We monitored the biasing with a single Housekeeping Board. The boards communicated with other electronics systems through the Bias-Board Backplane. A weatherproof Electronics Enclosure housed and protected the Bias system (the Readout system shared this enclosure).

MMIC Boards provided voltage and current to power the MMIC amplifiers in the modules. Each MMIC Board could bias two MABs (14 modules); there were seven MMIC Boards in total. The boards provided voltage sources for the MMIC gates and drains. We stabilized each drain-voltage source with an op-amp feedback loop ([Seiffert, 2008](#)). A 10-bit (2^{10} settings) digital-to-analog converter (DAC) controlled each voltage. (The DACs were Linear Technologies LTC1660 chips. Each LTC1660 has eight independent DACs. We used an “address” 1–8 to distinguish different DACs on the same chip. The Preamp and Phase-switch Boards used the same chips. Unless otherwise noted, “DAC” means one of the DACs controlling a bias-board output.) Using a DAC to control each bias setting allowed us to reliably save and reapply the biasing state of the entire receiver. Moreover, with digital control we could optimize the performance upon adjusting the settings ([§1.3.3](#)).

Phase-switch Boards provided control to the PiN diodes in the module phase switches. Each Phase-switch Board could bias three MABs (21 modules); there were five Phase-switch Boards in total. For each phase switch, the Board forward biased one PiN diode and reverse biased the other diode. The reverse bias was always -1.8 V ⁵. We adjusted the forward bias (controlled by a DAC) to make the phase-switch transmission nearly independent of the bias setting (Figure 2.6 and [§1.3.3](#)). Because different diodes had different transmission in

5. §2.5.3 of [Brizius \(2011\)](#)

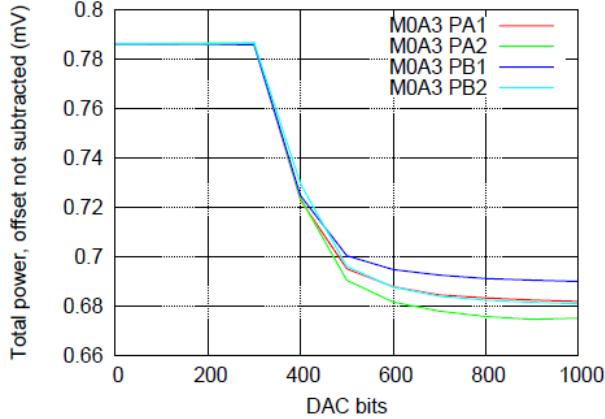


Figure 2.6: The phase-switch diodes required forward bias (positive DAC setting) to transmit power. Lower voltage corresponds to higher transmission power. The four phase-switch diodes (different colors) in each module had different transmission plateaus. We biased all phase-switch diodes in the plateau region and used double demodulation to remove the systematic effects of the transmission differences. SOURCE: [Tajima \(2009d\)](#)

this plateau region, we could not balance the transmission between the two phase-switch states⁶. Instead, we used double demodulation to remove the $I \rightarrow Q/U$ leakage that would otherwise be caused by the imbalance ([QUIET Collaboration, 2012b](#)). At 4 kHz (50 Hz) the Phase-switch Board switched the bias direction of both diodes in the L (R) phase switch. The Readout system provided synchronized clocks to synchronize the switching of the two switches. Because the AIB filtered the phase-switch bias, the switching transition was not instantaneous. We made the Readout system reject 17.5 μ s of data during each transition when the signal was unstable (Figure 2.7).

Preamp Boards biased and read out the detector diodes. Each Preamp Board could bias two MABs (14 modules); there were seven Preamp Boards in total. The bias circuit for each diode combined a floating⁷ current source with a differential voltage amplifier⁸. A DAC controlled the bias current. The amplifier circuit subtracted a (DAC-adjustable) voltage

6. We did balance the transmission in Q band.

7. i.e. neither side of the diode was at the module ground

8. §4.4.3 of [Bischoff \(2010\)](#) and [Bischoff \(2006b\)](#)

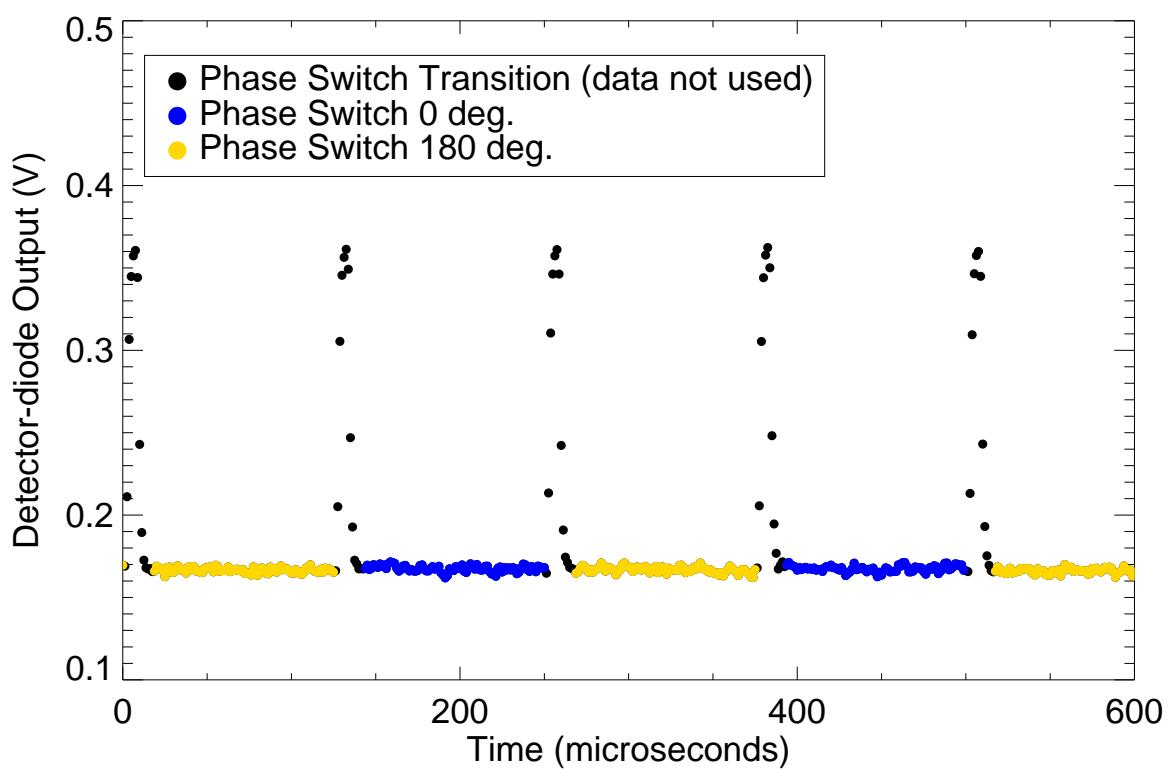


Figure 2.7: A phase switch modulated the sign of polarization at 4kHz. We discarded the data during the phase-switch transition.

offset so that the diode voltage, output by the Preamp Board, was centered in the dynamic range of the Readout system⁹. §2.3 provides details about the amplification and filtering applied by the Preamp Board.

One Housekeeping Board monitored all the MMIC and phase-switch biases. (We monitored the diode bias by recording the DAC settings and diode voltages.) We added several current sense resistors and voltage-probe points to each bias board. Multiplexers on the bias boards selected one of these measurements to send to the Housekeeping Board. The Housekeeping Board selected the measurement from one bias board or a thermometer (recording the temperature in the cryostat or Electronics Enclosure) to digitize with a 16-bit analog-digital converter (ADC). We monitored 1774 biases and temperatures (Appendix D), and the multiplexing rate was 500 Hz. Thus we measured each every 3.5 s. The multiplexing cycled only during phase-switch transitions (during which we already discarded the data) so the multiplexing transition did not affect the module signal. Analog opto-isolators separated the bias circuitry from the monitoring circuitry. The Readout and Data Management systems controlled this multiplexing.

The Bias-Board Backplane housed the bias boards and provided an interface to them. It was a modified VME (VERSAmodule Eurocard bus) 6U backplane with a partial metal enclosure that supported the bias boards (Figure 2.8). The backplane P2 (bottom) pins were straight-through and connected the AIB Cables to the bias boards (Bogdan, 2005). We connected the P1 (top) pins connect together so the bias boards could communicate. These pins provided a path for the bias boards to send a monitored voltage to the Housekeeping Board, a path for the housekeeping digital output to the Readout system, distribution of the housekeeping multiplex address, distribution of the phase-switch clocks, and digital control of the DACs. Table 2.2 lists the location of each board in the backplane.

The Electronics Enclosure housed and protected the Bias and Readout systems. It was an

9. Because of Type-B glitching (§2.3) the DD data acquired sensitivity to the voltage offset.

Table 2.1. Bias-board Locations

Slot	Board
1	Preamp 1
2	Preamp 2
3	Preamp 3
4	Preamp 4
5	Preamp 5
6	Preamp 6
7	Preamp 7
8	thermometer for Electronics Enclosure regulation
9	Housekeeping
10	MMIC 1
11	MMIC 2
12	MMIC 3
13	MMIC 4
14	MMIC 5
15	MMIC 6
16	MMIC 7
17	Phase-switch 1
18	Phase-switch 2
19	Phase-switch 3
20	Phase-switch 4
21	Phase-switch 5

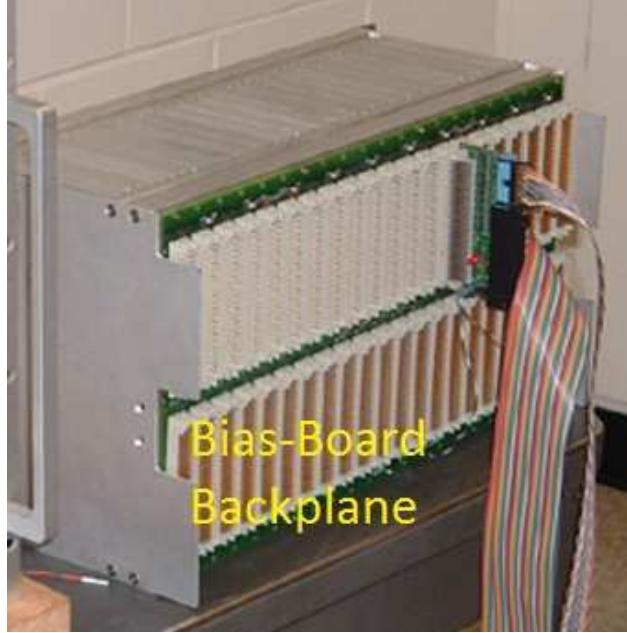


Figure 2.8: The bias boards entered through the face hidden in this photo. The AIB Cables connected to the lower set of connectors shown.

insulated 54" \times 25" \times 25" box we mounted next to the cryostat (DDB Unlimited OD-78DDXC, Figure 2.9 and §2.5.2 of [Brizius \(2011\)](#)). It contained the Bias-Board Backplane, power supplies, CPID, and Receiver PC¹⁰ (Personal Computer). Because the responsivity could vary with the electronics temperature¹¹, we regulated the temperature of the Electronics Enclosure.

2.2.1 Temperature Regulation

I designed a proportional-integral-derivative (PID) feedback-loop temperature-regulation system for the Electronics Enclosure (Figure 2.10 and [Buder \(2008\)](#)). The system consisted of a PID controller (Omega CNi16D22-EI, denoted “WPID”), thermometer (McMaster-Carr 6577T35 resistance temperature detector), solid state relays (SSR, Carlo RJ1A23A20U),

10. a rackmount computer that ran software for the Data Management system

11. We limited this effect to 0.3%/K ([Nagai, 2012e](#)). In Q band, the effect was larger due to a different MMIC Board design ([Bischoff, 2010](#); [Brizius, 2011](#)).

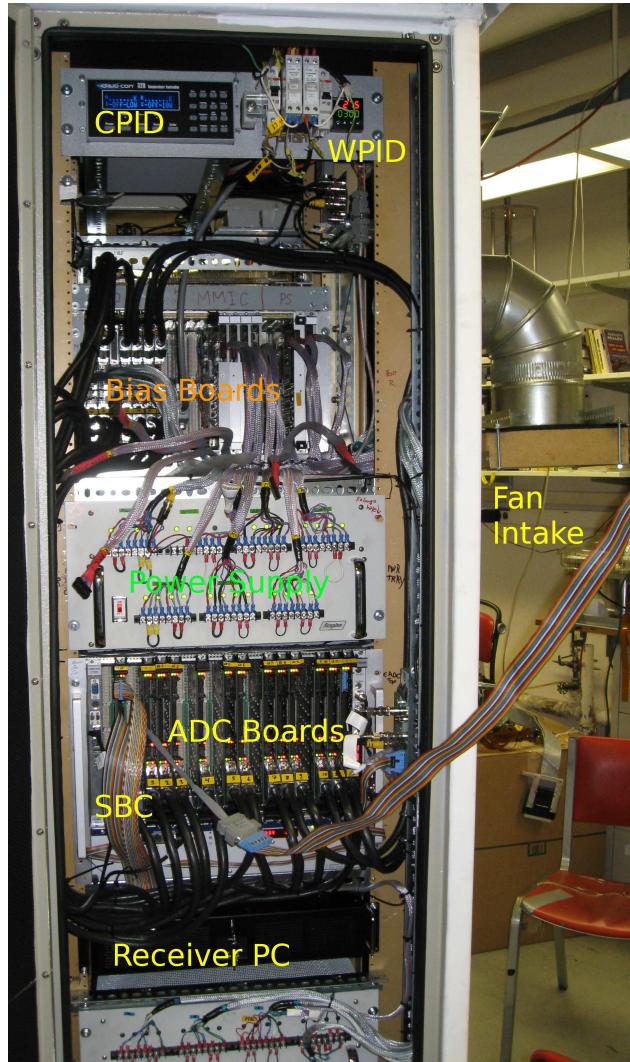


Figure 2.9: The Electronics Enclosure housed and protected the Bias and Readout systems.

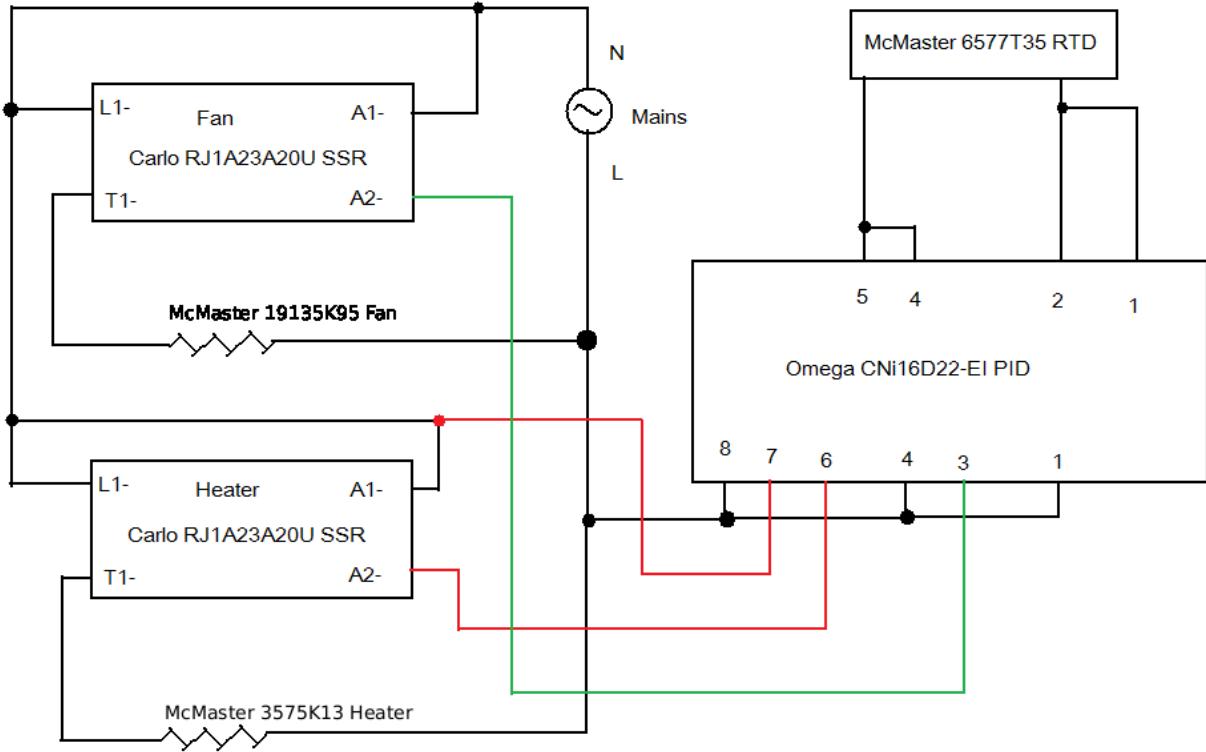


Figure 2.10: I made a proportional-integral-derivative (PID) system to regulate the Electronics Enclosure temperature. The numbers on the PID controller correspond to the physical numbered terminals. Wires of different colors do not connect except at filled circles. “L” represents the live electrical phase. When the PID turns on output 1, terminals 4 and 6 become connected. This connects the live phase to A2 on the Carlo SSR, activating the relay. The SSR connects L1 to T1, providing power to the heater. The same sequence applies to the fan (output 2) with terminals 1 and 3.

heater (McMaster-Carr 3575K13), and fan (McMaster-Carr 19135K95). I achieved good regulation performance ($< 1^{\circ}\text{C}$) on the time scale of a single observation (1 hour); however, the temperature changed significantly (20°C) throughout the season (§4.5.1). I made a monitoring circuit to let the Data Management system record the heater and fan activity.

The WPID controlled the Enclosure temperature by alternately activating the heater or using the fan to bring cold air into the Enclosure. The thermometer measured the temperature in the Bias-Board Backplane. Based on this temperature, the WPID activated one of the outputs to bring the temperature to the setpoint (35°C or 40°C depending on the

time of year, see §4.5.1). Each output activated an SSR that supplied power to the heater or fan. I mounted the heater directly above the Bias-Board Backplane; the heater supplied up to 1.1 kW. The fan brought cold ($-10\text{--}+15^\circ\text{C}$) air in above the Bias-Board Backplane. Air circulation fans (McMaster-Carr 1985K11) kept the air inside the Enclosure well-mixed. The WPID had an ethernet interface, and I wrote a custom control program¹² for remote control and automated operation by the Data Management system.

2.2.2 WPID Output Monitoring Board

The WPID did not report when it was turning on the fan and heater; therefore, I built an auxiliary monitoring circuit (“WPID Output Monitoring Board,” Figure 2.11). Linear power supplies converted convert the AC fan and heater voltages to 5 V DC. The WPID Output Monitoring Board supplied this voltage to the Receiver PC parallel port. A buffer (ON Semiconductor MC74VHC1GT125DT1G) and voltage divider (resistances 470Ω and 560Ω) protected the port from overvoltage damage. I made the Data Management system sample the parallel-port input to record the heater and fan activity (§3.1.2).

2.3 Readout

The Readout system amplified, filtered, and demodulated the signal. Preamp Boards provided analog amplification and low-pass filtering. ADC Boards digitized the module signals and demodulated them synchronously with the phase switching. A Time-code Reader provided the time and reference clocks for all electronics systems.

We used the Preamp Boards to measure the detector-diode voltage. Three analog amplifiers provided a gain of ≈ 130 ¹³. This amplifier chain contained low-pass filtering with three

12. Source code at <https://cmb.uchicago.edu/svn/ibuder/PIDControl>

13. §4.4.3 of Bischoff (2010). Late in the design stage we moved a factor of two in gain from the ADC Board to the Preamp Board (Bischoff, 2006a). As a result, some QUIET documents incorrectly give the Preamp Board gain of 65.

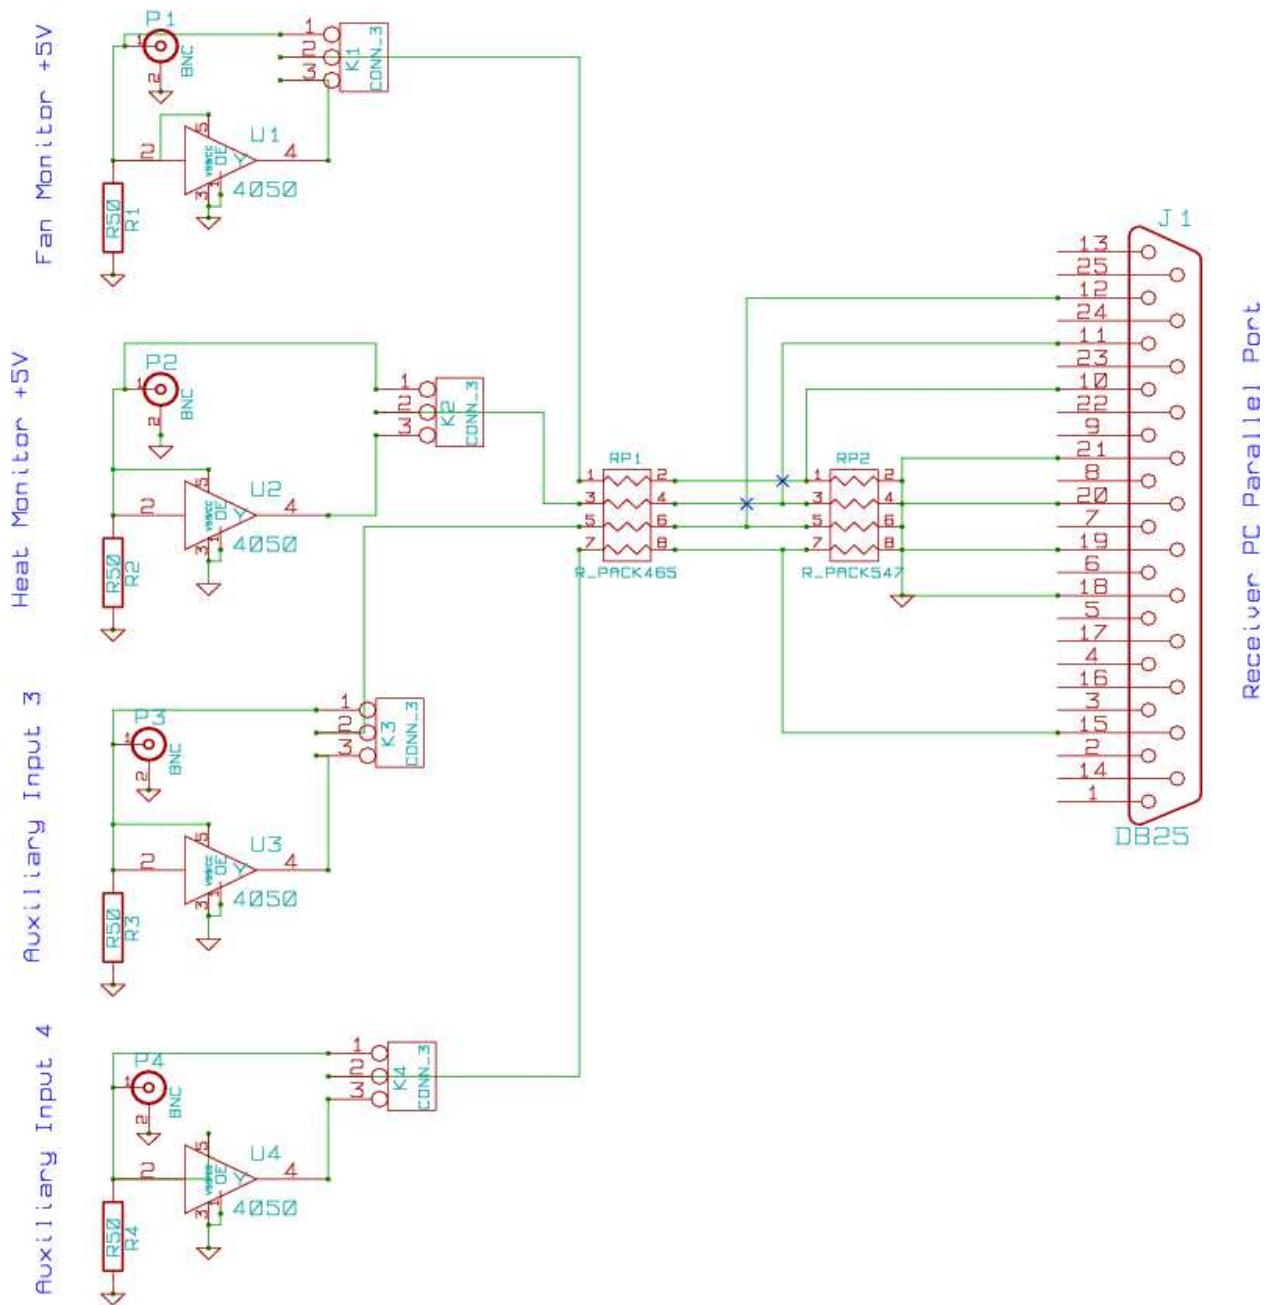


Figure 2.11: The WPID Output Monitoring Board sent the WPID heater and fan activity to the Data Management system for recording.

poles at 3-dB frequencies of 160 kHz, 340 kHz, and 32 MHz. This filtering prevented aliasing with the 800-kHz digitization sampling. As described in §2.2, the board subtracted a voltage offset to compensate for the diode-bias voltage. The Preamp Board output stage sent the amplified, filtered voltage to ADC Boards using one pair of wires (differential signaling) per diode.

ADC Boards digitized the diode voltages. Each ADC Board had 32 channels; however, we left four channels unused so each board digitized 28 diodes (seven modules, one MAB). Each channel had a low-pass filter (3 dB at 600 kHz) before digitization by 18-bit Analog Devices AD7674 Successive Approximation Register ADCs. The ADCs sampled at 800 kHz. Each ADC Board had a field-programmable gate array (FPGA), which accumulated the samples from the 32 ADCs on that Board. The FPGA on one ADC Board, designated the “Master ADC Board¹⁴,” generated the 4-kHz and 50-Hz signals used by the Bias system to control the phase-switch modulation. We distributed these signals to all ADC Boards, and the FPGA on each ADC Board used them to demodulate the detector-diode signals synchronously with the phase-switch modulation.

Normally we recorded three 100-Hz data streams for each diode. We recorded short (1024-sample) “snapshots” of 800-kHz data for diagnostic and monitoring purposes; however, there was insufficient storage to record at that rate continuously. To create the TP data stream (§1.3.1), the FPGA summed all 800-kHz samples, regardless of phase-switch modulation. The TP stream was sensitive to Stokes I , and we used it for calibration and monitoring. To create the DE data stream, the FPGA differenced samples, assigning different signs to the two 4-kHz phase-switch states. We differenced two adjacent DE samples in post-processing to demodulate the 50-Hz phase state switching and create the DD stream. In the “quadrature” stream, we differenced 4-kHz samples 90° out of phase with the 4-kHz switching (Figure 2.12). The quadrature data had the same noise as DE data and no signal. We used quadrature

14. The Master ADC Board did not accept signals from the modules. Therefore, the total number of ADC Boards was 14.

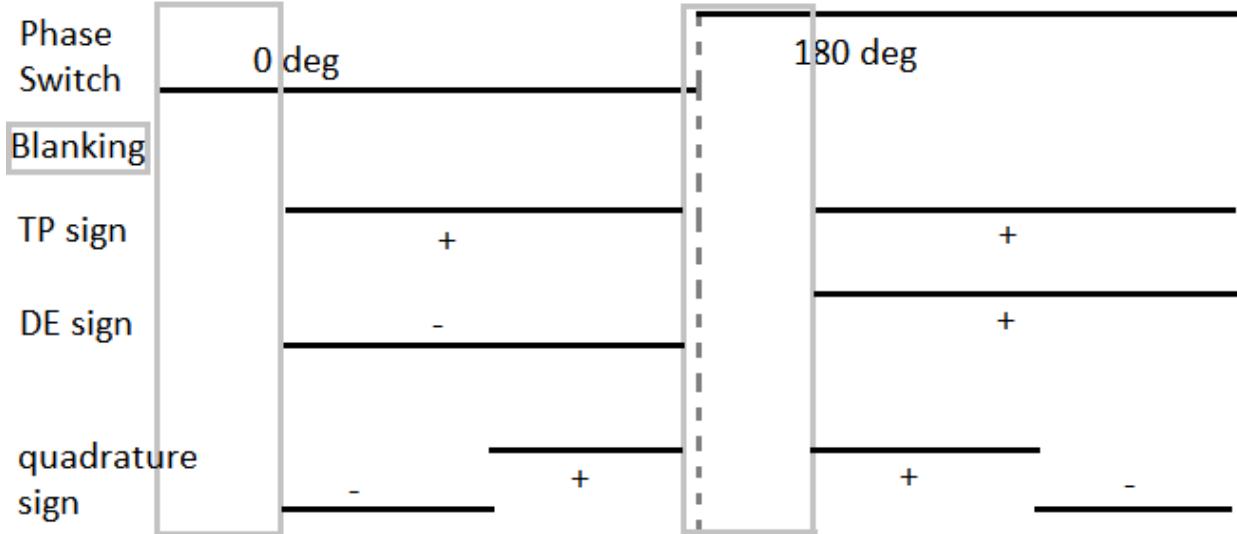


Figure 2.12: We implemented digital demodulation on the ADC Boards. We discarded data during phase-switch transitions (“Blanking”). The TP stream summed all data regardless of the phase-switch state. The DE stream differenced data between the (two 4-kHz) phase-switch states. Quadrature data was 90° out of phase so that we canceled the signal. (I exaggerated the fraction of data blanked, 14%, for ease of viewing.)

data to monitor potential contamination and understand the detector noise properties. The three streams (TP, DE, and quadrature) together constituted the “radiometer data.” As mentioned in §2.2, the FPGA did not accumulate data in any stream when the phase-switch state was changing. This blanking rejected 14 out of every 100 samples, reducing the effective integration time.

The ADC Boards had a small differential non-linearity in their response. At intervals of 1024 bits¹⁵, the ADC output had a jump discontinuity of size < 40 bits (Figure 2.13). The size was negligible compared to the large 1/f noise in the TP data. However, the effect in DE data was not negligible. When the average 800-kHz voltage level was near a discontinuity, the voltage difference between the two phase-switch states caused the level to move across the discontinuity. This caused the discontinuity to be added to the DE data. The 800-kHz noise (much higher per sample than per 100-Hz sample) smeared out the discontinuity, reducing its

15. There was one exceptional channel that had a different interval.

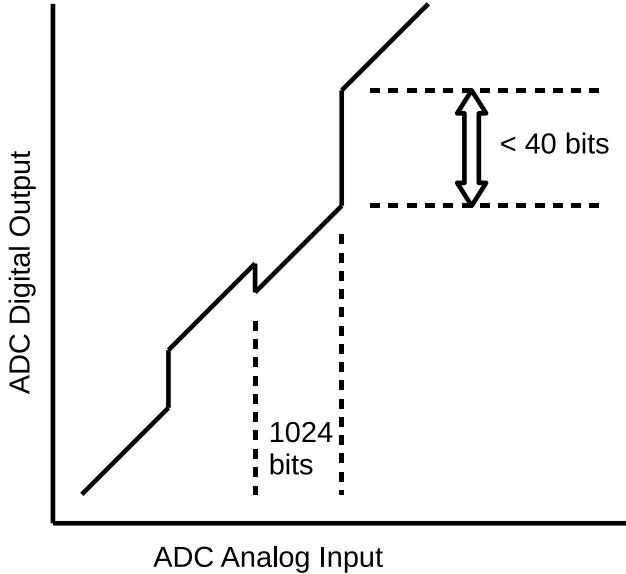


Figure 2.13: Type-B glitches were jump discontinuities (differential non-linearity) in the ADC response. The jumps were evenly spaced every 1024 bits; however, the jump size and direction were different for different jumps. When the input was near a jump, the difference between the two phase-switch states caused the jump size to add to the DE data.

amplitude but also increasing its width in bit space (the jump can occur whenever the signal level plus noise encounters the discontinuity). We called this effect “Type-B glitching¹⁶.” We mitigated the effects by using the Preamp Board offsets to move the typical voltage levels away from the discontinuities. However, some effect remained, and I corrected for Type-B in analysis (§5.1.5).

The ADC Backplane contained the ADC Boards, Time-code Reader, and a single-board computer (SBC, part of the Data Management system). The ADC Backplane was a Wiener Series 6000 VME crate; we put it in the Electronics Enclosure. The Master ADC Board used the VME interface to distribute the phase-switch control signals (and all other clock signals for synchronization) to the other ADC Boards.

The timing hardware consisted of a Time-code Reader and Auxiliary Timing Board

¹⁶. We believe Type-B was caused by a combination of bad ADC clock signals and bad ADC chips. Type-B disappeared when we operated the ADCs at 400 kHz. Some ADC channels did not have this glitching, and the jump sizes for the channels that did glitch were not identical. However, the exact mechanism in the ADC causing Type-B is not known.

(ATB). The Time-code Reader synchronized the electronics to an external Global Positioning System (GPS) derived time. The ATB distributed timing signals from the Time-code Reader to the Master ADC Board, which synchronized the other ADC Boards.

The Time-code Reader was a Symmetricom TTM635VME-OCXO ([Symmetricom Inc., 2005](#)). We used an IRIG-B (Inter-range Instrumentation Group Mod B) amplitude-modulated time code from the GPS receiver in the control room (§4.1) as the reference. From the time code we created synchronized 1-Hz and 10-MHz clock signals. The ATB distributed these clocks to the Master ADC Board, which used them to synchronize all the ADC Boards¹⁷.

We made the ATB to convert logic levels between the Time-code Reader and Master ADC Board. The Time-code Reader output levels were TTL/CMOS, but the Master ADC Board could only receive low-voltage differential signals (LVDS). We designed a small circuit board to convert between them (Figure 2.14). The board used National Semiconductor DS90C031 LVDS Differential Line Drivers.

17. §2.5.4 of [Brizius \(2011\)](#)

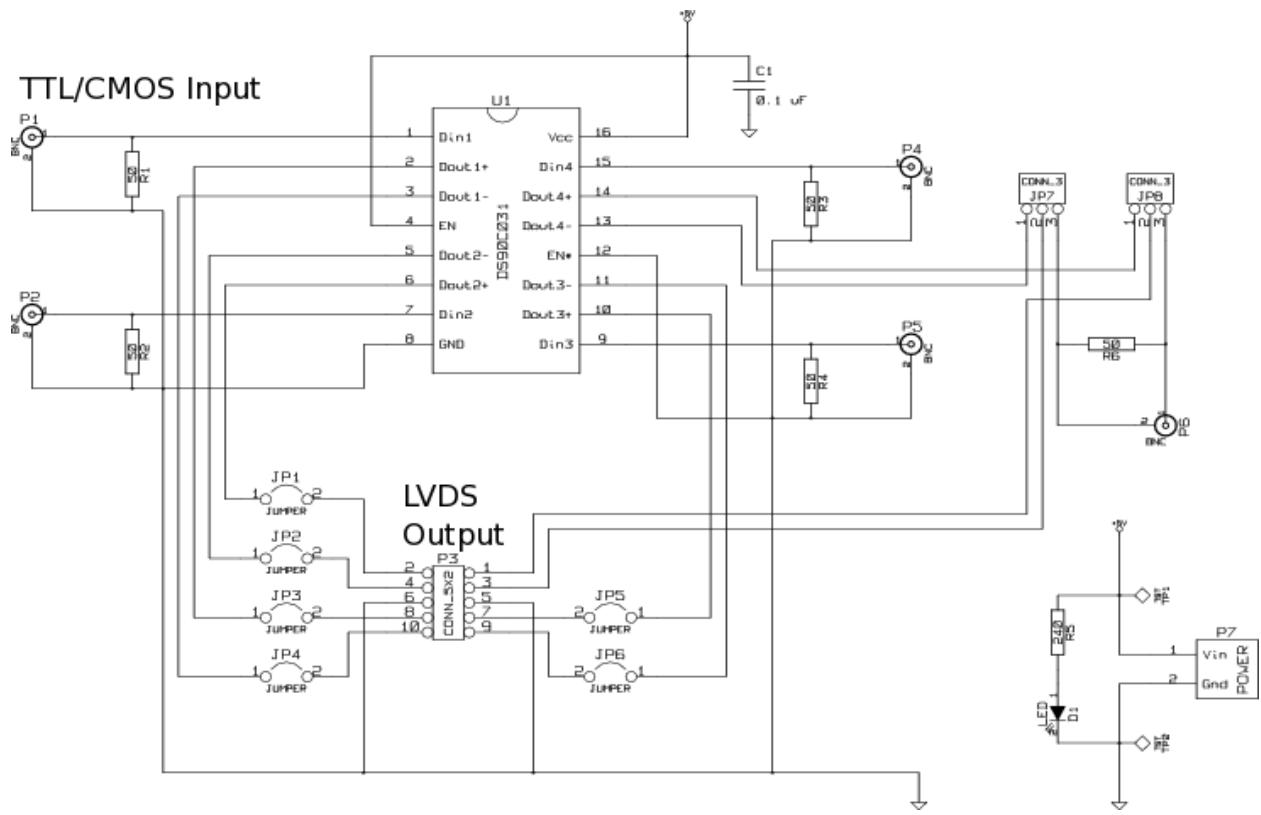


Figure 2.14: The Auxiliary Timing Board (ATB) converted logic levels between the Time-code Reader and Master ADC Board.

Chapter 3

Data Management

The first 90 percent of the code accounts for the first 90 percent of the development time. The remaining 10 percent of the code accounts for the other 90 percent of the development time.

Tom Cargill

The Data Management system was responsible for sending commands to the Bias system and mount pointing control system, receiving data from the Readout system, synchronizing data from different readout subsystems, and recording the data. We created several software processes that ran on several computers to accomplish these main functions. We organized the processes in a hierarchy. For example, the lowest level bias control process (`adc_server`) understands and implements only very simple commands e.g. “change the bias output of DAC #N to m bits.” Whereas the highest level control process (`Receiver Control Panel`) has a command “change the bias so the receiver is on.” Each process translated high level commands to low level commands (or low level data to high level data) appropriate for the other processes it communicated with. The subsections of this chapter detail these interactions.

We divided software processes into two broad categories: Online Software and Control Software. Online Software processes were always running and rarely used by the observers directly. Control Software processes had user-friendly graphical interfaces. The observers used them to schedule and start observations or monitor the system status. Two additional major Data Management components were the data files, whose format is described in §3.4,

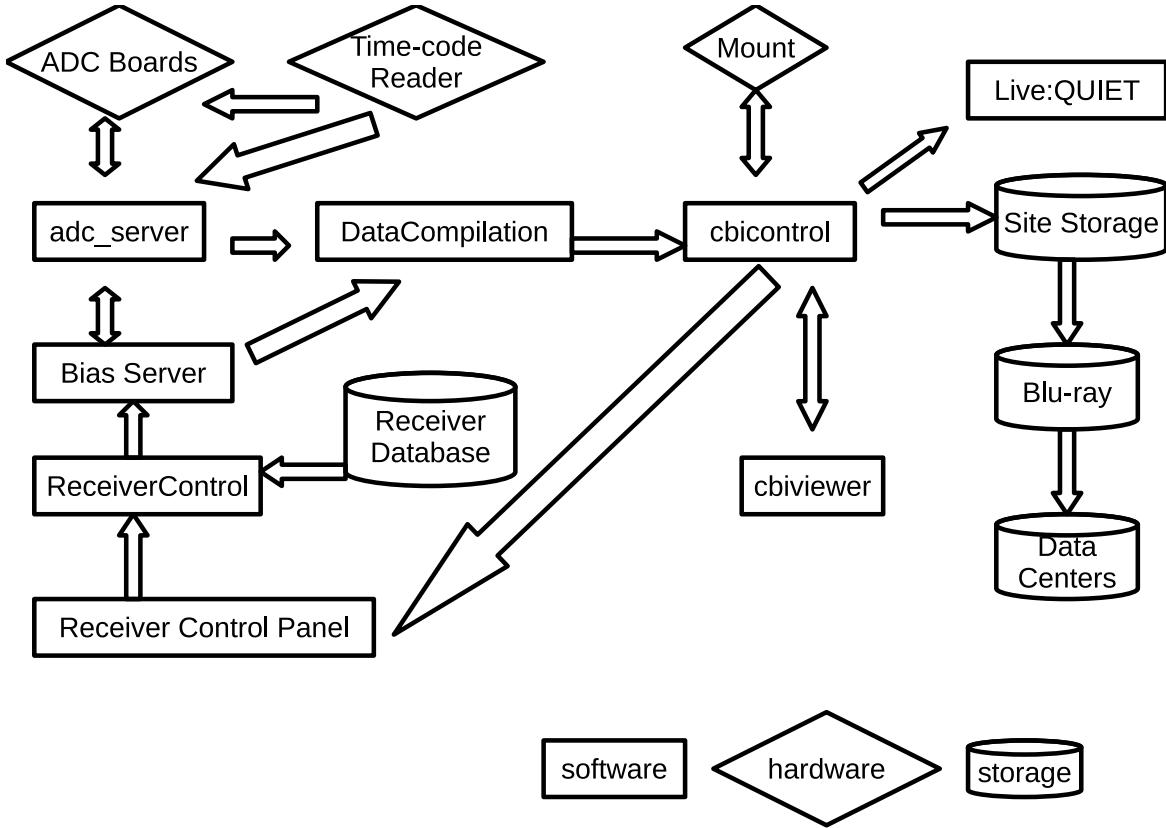


Figure 3.1: The Data Management system recorded, distributed, and stored the data. I wrote software processes that were its major components.

and Receiver Database. The Receiver Database contained configuration information, e.g. which electronics board was connected to which module, which was used when translating command levels. Figure 3.1 is a simplified diagram of the Data Management system.

3.1 Online Software

The Online Software was responsible for the continuous readout and recording of the data. In addition, we periodically (or sporadically) generated bias commands. The Online Software translated these commands to the lowest level and sent them to electronics hardware. The Online Software helped us monitor the DAQ and command implementation status and

automatically detected and reported several types of problems (Appendix A) via the Control Software.

We created five processes to accomplish these functions. The `adc_server` ran on the single-board computer (SBC, a GE Fanuc VME-7700) in the ADC Backplane and was responsible for direct communication with electronics hardware or firmware. The `DataCompilation` process ran on the Receiver PC in the Electronics Enclosure and was responsible for combining data from different readout systems and distributing them to recording and monitoring processes. The `ReceiverControl` process ran on the Receiver PC and allowed us to enter high-level bias commands. The `Bias Server` process, which translated electronics bias commands, and `PeripheralServer`, which interfaced with auxiliary pressure sensors, temperature controllers, etc. are described in [Bischoff \(2010\)](#).

3.1.1 `adc_server`

The process `adc_server`¹ was the lowest level of the Online Software. It communicated directly with the ADC Boards and Time-code Reader using the VME64 interface on the ADC Backplane. The process `adc_server` received commands from and sent data to higher-level Online Software processes via 100BaseTX ethernet ([GE Fanuc Embedded Systems, 2007](#)). The process `adc_server` was the only user process running on the SBC, and it ran continuously as long as the ADC Backplane was powered. I divided the functions among software threads. Since the SBC had only one Central Processing Unit (CPU) core, only one thread was running at any given time. The operating system ([Vanderlinde, 2007](#)) could switch rapidly between them, but some synchronization was necessary. Table 3.1 provides an overview of the threads and their functions.

The `msg_serve` thread waited for a network socket connection on port 5005. After a

1. Source code is available at https://cmb.uchicago.edu/svn/kapner/quiet/crate_fs/branches/quadrature/root/vme_adc_threaded/src

Table 3.1. Summary of `adc_server` Threads

Thread Name	Function	Priority
<code>msg_server</code>	Send warning and error messages to the network for recording	8
<code>fourhz_write</code>	Read data from ADC Boards, read time from Time-code Reader	10
<code>data_server</code>	Send radiometer data to <code>DataCompilation</code>	8
<code>data</code>	Send snapshots to <code>DataCompilation</code>	8 (9 while sending a snapshot)
<code>bias</code>	Implement commands from <code>BiasServer</code> , send biasing status to <code>BiasServer</code>	4
<code>hk_server</code>	Send housekeeping data to <code>DataCompilation</code>	8

connection was established, it sent the `adc_server` software version (4 bytes) and ADC Board firmware version (4 bytes) over the network. (Unless otherwise specified, all network communication began with sending these version numbers.) The thread then slept until there was a message to send. Other threads put error, warning, and informational messages into a buffer which can hold 2000 such messages for sending. I incremented a semaphore when a new message was put into the buffer. The thread checked for an acknowledgment that the message was received before deleting it from the buffer. If there was no acknowledgment, the thread closed the socket connection and waited for a new one.

The `fourhz_write` thread waited for the Master ADC Board to signal that a new data frame was ready to be read. It did this by setting the 6th bit² of the word at VME memory address 0x0³ to 1. When a frame was ready to be read, the thread first latched the time from the Time-code Reader. The thread then read the clock status from the Master ADC Board, radiometer data from each ADC Board, housekeeping data from the Master ADC Board,

2. i.e. the 2^5 place

3. All VME memory addresses have an offset that depends on the slot the ADC Board occupies. I suppressed these offsets here. See §5 of [Vanderlinde \(2007\)](#) for more details.

and the frame counter from each ADC Board. The thread checked that the frame counter was the same for all ADC Boards. This ensured that data from different ADC Boards were acquired at the same time.

I combined the data and status information in packets shown in Figure 3.2. The first 8 bytes of the packet were the timestamp from the Time-code Reader. The timestamp was a Unix time; the first 4 bytes were the integer number of seconds since 00:00:00 Coordinated Universal Time (UTC), January 1, 1970; the last 4 bytes were the integer number of microseconds⁴. Following the timestamp in each packet were the data. If the packet contained housekeeping data, these were 125 housekeeping multiplex addresses (16-bit integers) and values (16-bit integers). If the packet contained radiometer data, these were 25 samples (32-bit integers) each for 32 ADC Board channels for TP, DE, and quadrature per ADC Board. Finally, each packet ended with 16 bytes of status information (Appendix B.1). I temporarily stored the packets in a ring buffer capable of holding 60 s of data. After storing each packet, the `fourhz_write` thread incremented a semaphore to notify the `data_serve` or `hk_serve` threads that the packet was ready to be sent. If network problems or a crash in `DataCompilation` caused the ring buffers to become full, data was lost. In use, a problem that lasted for 60 s typically required human intervention to solve so the buffer size was not a limiting factor causing data loss.

Every 100 s the thread used the Time-code Reader time to update the SBC system time. This procedure maintained the synchronization between them to within 5 ms. I only used the SBC time to timestamp monitoring data e.g. snapshots, logs, and 1-Hz housekeeping data. Thus a possible 5-ms error in this time was acceptable.

The `data_serve` thread sent radiometer data packets over the network. It waited for a connection on port 5000. Upon establishing a connection, it sent the number of ADC

4. The use of Unix time here prevented proper representation of the time in the presence of leap seconds. No QUIET observations contained a leap second.

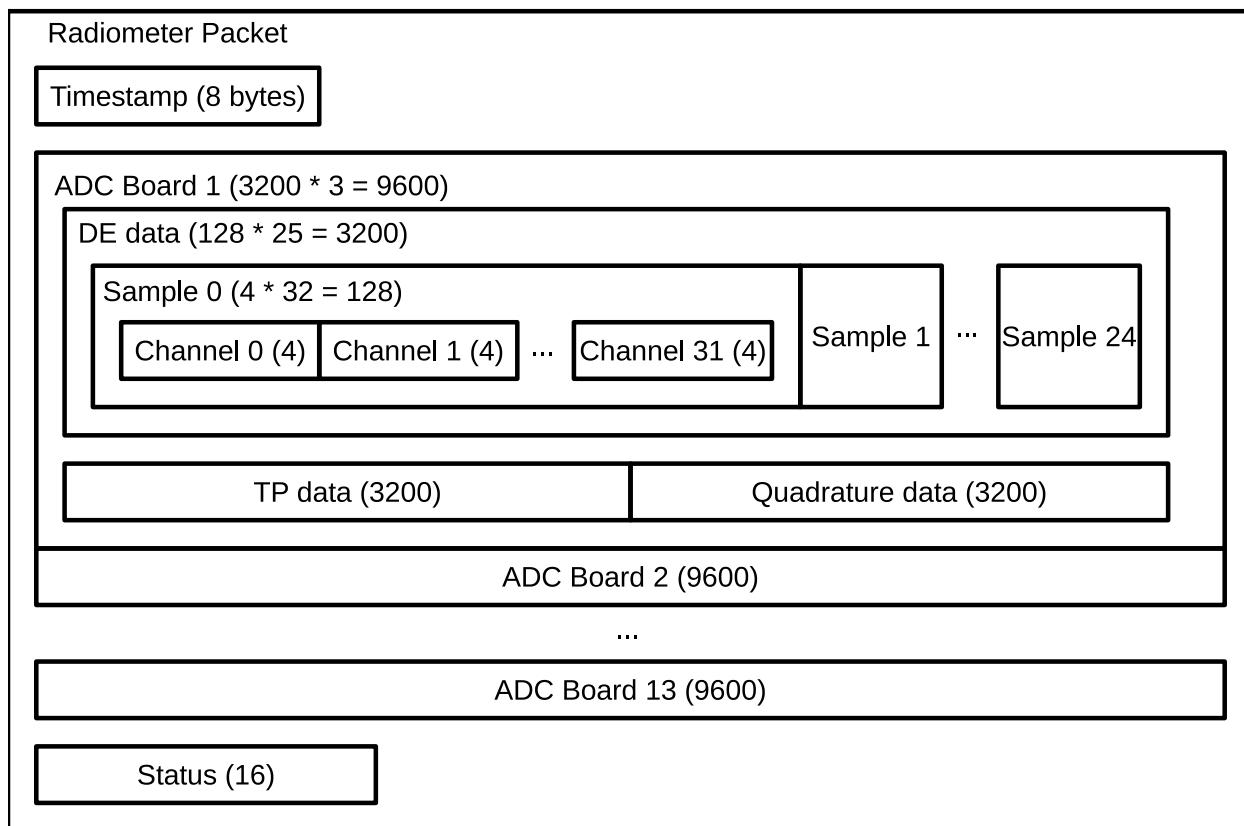


Figure 3.2: I made `adc_server` prepare packets containing 0.25 s of TP, DE, and quadrature data (collectively, “radiometer data”). Numbers in parentheses represent bytes.

Boards recording data⁵. Then it waited for a radiometer data packet to become ready to send. When a packet was ready, it sent it over the network socket. Then it waited for an acknowledgment from the other end of the connection. If the acknowledgment byte was ‘c’, then the thread erased the sent packet from the buffer and returned to waiting for a packet to be ready. If the acknowledgment byte was ‘q’, then the thread erased the packet, closed the network connection, and returned to waiting for a connection. If no acknowledgment was received, then the thread closed the connection, but did not erase the packet from the buffer. This action ensured that data was not removed from the buffer before the data was properly stored in `DataCompilation`.

The `hk_serve` thread sent housekeeping data packets over the network. It waited for a connection on port 5003 and then waited for a housekeeping data packet to become ready to send. I used the same packet sending and acknowledgment system as in the `data_serve` thread.

The `data` thread acquired and sent snapshots. It waited for a connection on port 5001. Upon establishing a connection, it sent the number of ADC Boards recording data. Then it wrote the value 0 to the 4-byte word at VME memory address 0x30 on the Master ADC. This instructed the ADC Boards to take a snapshot. The thread waited 10 ms for the instruction to be processed and then wrote the value 1 to the same VME memory address. This instructed the ADC Boards to save the snapshot. The thread waited 1 ms for the instruction to be processed. The thread then stored the SBC time to be used as the timestamp for the snapshot. The thread then sent a packet over the network consisting of:

- The 8-byte timestamp (see above for format).
- The 2^{17} -byte snapshot data for each ADC Board.

5. In W band the Master ADC Board did not acquire radiometer data; thus the number of ADC Boards recording data was the number of Slave ADC Boards. In Q band, all ADC Boards (Master or Slave) recorded data.

The thread then waited 1 ms before writing 0 to the VME memory address, instructing the ADC Boards to release the saved snapshot. The thread finally closed the connection and returned to waiting for a connection.

The `bias` thread implemented commands from `BiasServer` and reported the status of module biasing. It waited for a connection on port 5002. After establishing a connection, the action depended on the request received (Appendix [B.2](#)).

Since only one `adc_server` thread can run simultaneously, I assigned priority and implemented synchronization among them. My goal was to ensure that the `fourhz_write` thread always read data from the ADC Boards without delay and assigned the proper timestamp while minimizing delay in the tasks of other threads. The `adc_server` threads had static scheduling priority; I broke ties within a priority level by round-robin scheduling. Since the `fourhz_write` thread had the most critical operation, I assigned it the highest priority. The `bias` thread had the lowest priority because several bias commands take a long time (10s or more), and the other threads should not be blocked while such long commands occur. The other threads had equal priority except that the `data` thread had higher priority while responding to snapshot requests. Because the snapshot timestamp was assigned by the thread, it should not be interrupted to provide the most accurate timing. However, the `fourhz_write` thread will still interrupt snapshot requests if necessary. When idle or not in a time-critical operation, I made each thread sleep or yield to allow other threads to run.

The threads used synchronization techniques to share resources. I protected communication with the ADC Boards by a mutex⁶ so that they need only keep track of one request at a time. The `fourhz_write` thread used a semaphore to indicate when new data were ready to be sent. I used a producer-consumer model to ensure that ring buffers were not accessed unless they were ready to be read or written. I protected buffers that had multiple producers (e.g. the message buffer) by a mutex.

6. A mutex ensures that multiple threads do not access a shared resource (memory, VME communication, etc.) simultaneously ([IEEE and The Open Group, 2012b](#)).

I conducted tests to verify that timestamps assigned by `adc_server` were accurate. I checked the difference between the Time-code Reader timestamp available over VME and the time implied by its 1-Hz output was $< 35 \mu\text{s}$. This difference included the communication delay between the Time-code Reader and `adc_server`. To check for delays between the Time-code Reader and ADC Boards, we injected the Time-code Reader 1-Hz clock pulse into an ADC channel and found the pulse in the recorded radiometer data. This test confirmed a 20-ms delay (understood as ADC firmware buffering) also found from Moon observation (QUIET Chicago Group, 2008). To check that timestamps would be regular, I ran the DAQ system uninterrupted for 20 hours. All timestamps were within 1 ms of the periodic expectation. The Moon observations limited any overall offset in the timing system to $< 1 \text{ ms}$ (Kusaka, 2010b).

I confirmed that data were not lost between acquisition and long-term storage. The frame counter (see above) was a unique identifier for each data packet. I ran the system for 20 hours, and there were no missing packets. This test also checked possible loss in `DataCompilation`. To check that the firmware was assigning the frame counter correctly, I confirmed the timestamp incremented 250 ms for each frame counter increment, as expected. I also confirmed that packets were not duplicated⁷. In shorter tests, I introduced simulated network problems (e.g. temporarily disconnecting an ethernet cable) and confirmed that the system recovered as long as the problem duration was shorter than the data buffer length.

3.1.2 DataCompilation

The process `DataCompilation`⁸ combined multiple data sources into a single stream for later storage. It ran on the Receiver PC continuously during data taking. I wrote it as a

7. Duplication was a problem during software development. Although duplicated packets could be removed in post-processing, to keep the post-processing software simple I eliminated duplicates in the Online Software.

8. Source code is available at <https://cmb.uchicago.edu/svn/ibuder/RCS>

Table 3.2. Summary of DataCompilation Threads

Thread Name	Function
demod	Receive radiometer data from <code>adc_server</code>
hk	Receive housekeeping data from <code>adc_server</code>
data_server	Send system state information and recent data to other Online Software processes
bias	Receive bias state information from <code>BiasServer</code>
archive_aggregator	Combine different data sources into a single packet
timestamp	Timestamp the <code>DataCompilation</code> log file every second
snap	Request a snapshot from <code>adc_server</code> every hour and write it to disk
per	Receive auxiliary housekeeping data from <code>PeripheralServer</code>
wpid_output	Record WPID heater and fan output information
pdu_read	Record auxiliary housekeeping data from Electronics Enclosure smart power strip
data_arc_service	Send combined packets to <code>DataArc</code>
cbi_service	Send combined packets to <code>cbicontrol</code>

multi-threaded C++ program. The threads and their functions are summarized in Table 3.2.

The `demod` thread received radiometer data from `adc_server` and performed preliminary processing on it. It created and maintained a socket connection to the `adc_server` `data_server` thread. Upon receiving a packet, it first checked if the frame counter was the same as the most recently received frame counter; if they were the same, the newly received frame was a duplicate and discarded. Otherwise, the thread unpacked state information the firmware has stored in the sign bit of the DE data (Vanderlinde, 2007). The sign bits of channels 0–11 on each ADC Board comprised the number of samples accumulated when the 4-kHz clock was in the up state. (Channel 0 contained the 1s place, channel 1 the 2s place, channel 2 the 4s place, etc. in the base-2 representation.) The sign bits of the quadrature data of channels 0–11 comprised the number of samples accumulated when the quadrature sign was positive. Channels 16–27 contained the same information for the down state. Chan-

nels 30 and 31 contained the 1-Hz clock and 50-Hz clock, respectively. After extracting these data, the thread replaced the sign bit with the next most significant bit so that the value was in the standard two's complement representation. After this processing, the thread stored the resulting frame (250 ms of data and metadata) in a global buffer that can be accessed by other threads. It also put a copy in a queue (60 s long) for inclusion in the archive stream.

The `hk` thread received housekeeping data from `adc_server` and performed preliminary processing on it. I used the same algorithms as the `demod` thread except for the sign bit processing, which was not applicable to housekeeping data.

The `data_server` thread allowed other Online Software processes to receive the latest data over the network. It listened for a network connection on port 4001⁹. When it established a connection, it created a new thread (`data_serve`) to handle it. In this way, there can be up to five simultaneous connections to this interface. Each remote process may send a 4-byte request. The thread's action depended on the request (Appendix B.3). Table 3.3 summarizes the possible requests.

The `bias` thread connected to `BiasServer` on port 4002 and received the latest bias state (see `DATA_BIAS` in Appendix B.3) twice per second.

The `archive_aggregator` thread combined data collected by other threads into a single stream. This stream contained one frame per second. In each frame I combined radiometer, housekeeping (from all sources listed above), and biasing data; Online Software status flags; and the most recent snapshot number. The thread waited for four (the number needed to fill a single archive frame) radiometer and housekeeping frames to be collected by the `demod` and `hk` threads. I then aligned them according to a timing algorithm (see §3.2). Then I copied the latest biasing data, `PeripheralServer` data, WPID outputs, Electronics Enclosure monitoring, and snapshot number to the archive frame being constructed. I recorded the current version of the Receiver Database in the frame. I created a new frame number

9. Online Software port configuration is in `RCSports.h` in the RCS repository.

Table 3.3. Summary of `DataCompilation data_server` Requests

Name	Description	Value
DATA_STATUS	Send <code>DataCompilation</code> status flags	0
DATA_NADC	Send number of ADC Boards with data	1
DATA_DEMOD	Send latest radiometer data frame	3
DATA_SNAP	Send latest snapshot	4
DATA_HK	Send latest housekeeping data frame	5
DATA_BIAS	Send latest biasing state	6
DATA_PER	Send latest auxiliary housekeeping data acquired from <code>PeripheralServer</code>	8
DATA_NEWSNAP	Cause <code>DataCompilation</code> to request and acquire a new snapshot from <code>adc_server</code> , then send the identification number of the new snapshot	9
DATA_ENCLOSURE	Send monitoring data from Electronics Enclosure smart power strip	10
DATA_PID_OUTPUT1	Send output status of first WPID output	11
DATA_PID_OUTPUT2 ^a	Send output status of second WPID output	12

^aThere is support in `DataCompilation` for additional PID outputs, but they do not exist in hardware.

for the archive frame¹⁰. I also included the current Receiver PC time in the frame. This timestamp was not accurate enough¹¹ for data alignment or pointing reconstruction but was used for monitoring the Online Software status and debugging. When the frame was complete, the thread pushed it into two queues (60 s long) for transmission to **DataArc** and **cbicontrol**¹².

The **timestamp** thread wrote the current time (from the Receiver PC clock) to the log every second.

The **snap** thread requested and received snapshots from **adc_server**. It requested a new snapshot every hour or when signaled by a **data_serve** thread. The thread assigned each snapshot a number. I saved this number to disk on the Receiver PC so that snapshot numbers did not reset even if **DataCompilation** restarted. The thread wrote each snapshot to disk on the Receiver PC.

The **per** thread received the additional housekeeping data from **PeripheralServer**. The thread connected to **PeripheralServer** on port 9011 and requested a new data frame every second.

The **wpid_output** thread acquired the WPID output states. It sampled the Receiver PC parallel port at 100 Hz. The WPID Output Monitoring Board (§2.2.2) made the WPID output control signal available at the parallel port. The thread read the parallel port input (I/O port 0x0379 on the Receiver PC) with inb ([inb, 2012](#))¹³. The combination of the WPID

10. The archive frame number was four bytes so there were enough possibilities to avoid duplication during the entire lifetime of QUIET. However, the frame number reset to zero whenever **DataCompilation** restarted. Therefore, the archive frame numbers repeated. However, valid contiguous data cannot include a period when **DataCompilation** restarted. Thus, frame numbers were unique within a single scan.

11. I synchronized the Receiver PC time to the SBC time by Network Time Protocol, so it was indirectly derived from the same source as other timestamps. However, the synchronization was only guaranteed within a few seconds. Moreover, **DataCompilation** was not a real-time process so there can be unexpected delays in assigning the timestamp.

12. As a compile-time option, **DataCompilation** could be configured to send to either, neither, or both.

13. To access the parallel port in this manner, **DataCompilation** required root privileges. I typically compiled it to be executable by all users and setuid root. After granting access to the parallel port with ioperm ([iop, 2012](#)), **DataCompilation** dropped root privileges. Alternatively, **DataCompilation** may be run

Output Monitoring Board wiring and mapping of parallel port pins caused the WPID output 1 to be on bit 7 (2^7 s place) of the word read; output 2 was on bit 6. Each sample indicated whether the output was active at the time of sampling. The thread accumulated 100 samples into a frame and assigned each frame a timestamp based on the Receiver PC time. The WPID output timestamps should not be used for sub-second timing alignment.

The `pdu_read` thread collected the monitoring information from the Electronics Enclosure smart power strip. A separate process (`read_pdu.py`) connected to the power strip by ethernet ([CyberPower, 2012](#)) and wrote the resulting data to a file on the Receiver PC. The `DataCompilation` process read this file every second.

The `data_arc_service` thread sent archive frames to `DataArc`¹⁴. It waited for an archive frame to be put into the queue by `archive_aggregator` thread. Then it listened for a connection on port 4009. The next action depended on the message received from `DataArc` (Table 3.4 and Appendix B.4).

The `cbi_service` thread sent archive frames to `cbicontrol`. It waited for a frame to be put in the queue by the `archive_aggregator` thread. Then it waited for `cbicontrol` to be ready to receive a frame. Then it converted the frame into the format acceptable by `cbicontrol` ([Shepherd, 2008a](#)). I packed all data types as unsigned 4-byte integers. I reorganized the bias data to eliminate the unphysical addresses¹⁵. I reorganized the radiometer data so that the 100 samples from each channel were in consecutive order. Finally the thread sent the frame and waited for a new frame.

without root privileges; in this case, it did not read the WPID output information.

14. During normal W-band operation, we did not use `DataArc`. We only recorded the data with `DataArc` for lab tests and initial and special observations that bypassed `cbicontrol`.

15. See `BiasBitRepacker.h` in the RCS repository.

Table 3.4. Summary of `DataCompilation` and `DataArc` Communication

Message Name	Description	Value
AGG_NOMSG	Invalid message or communication error	0
AGG_FRAME	Message containing an archive frame (only sent by <code>DataCompilation</code>)	1
AGG_ACK	Message acknowledging receipt of frame (only send by <code>DataArc</code>)	2
AGG_ACK_REQ	Message requesting acknowledgment of receipt of frame (only sent by <code>DataCompilation</code>)	3
AGG_QUIT	The connection will be closed	4
AGG_NACK	The most recent frame was not received correctly (only sent by <code>DataArc</code>)	5
AGG_FLUSH_BUFFER	Instruct <code>DataCompilation</code> to flush the frame buffer	6

3.1.3 ReceiverControl

The process `ReceiverControl`¹⁶ allowed the Control Software and observers to manage the bias electronics. It ran on the Receiver PC whenever we changed the bias. It was a single-threaded C++ program and did not check for other instances. We took care to ensure that at most one copy was ever running. Otherwise, the bias state could be changed in an inconsistent and unexpected way. The behavior of `ReceiverControl` depended on command-line arguments (Table B.1. Further command details can be found by running `RC_functions_new ?`). The `ReceiverControl` process constructed lower-level commands which it sent to `BiasServer` and monitored the command results from `DataCompilation`.

The Receiver Database (Appendix D) allowed `ReceiverControl` to translate between bias identifiers and addresses. The database included the mappings between ADC channels, preamp channels, and module diodes; bias identifiers, cards, addresses, and DACs; house-keeping identifiers, addresses, and the corresponding bias boards; card locations, versions,

16. Source code is available at <https://cmb.uchicago.edu/svn/ibuder/RCS>

and serial numbers; module serial numbers, types, and locations; and housekeeping thermometer names, locations, types, and calibration constants. It also stored the opto-isolator conversion tables needed to convert housekeeping voltages to physical units. It also stored bias settings for rapid reapplication and reuse. The database stored its own version number for consistency.

3.1.4 DataArc

The `DataArc`¹⁷ process recorded the data to disk on the Receiver PC. We used it in the lab and for special observations at the site; normal site data were recorded by `cbicontrol`. I wrote it as a single-threaded C++ program. It connected to the `DataCompilation` `data_arc_service` thread and communicated as follows. First `DataArc` sent `AGG_FLUSH_BUFFER` to force `DataCompilation` to discard any old data. Then it waited for `DataCompilation` to send a message.

1. If the message was `AGG_FRAME`, `DataArc` wrote the frame to disk and sent `AGG_ACK`.
2. If the message was `AGG_ACK_REQ`, `DataArc` sent `AGG_ACK`.
3. In all other cases or if an error occurred, `DataArc` closed the connection and reestablished it.

3.2 Timing Alignment

I ensured that data were assigned correct timestamps and that data acquired from different sources at the same time were combined in the same frame. On the hardware side, the Time-code Reader provided absolute time from GPS and clocks that regulated the ADC. The ADC firmware ensured that frames were aligned with the Time-code Reader time (meaning that a

17. Source code is at <https://cmb.uchicago.edu/svn/ibuder/RCS>

frame began at 0.000 s) and that the interval between radiometer (housekeeping) samples was 10 ms (2 ms). On the software side, `adc_server` combined the ADC data with timestamps from the Time-code Reader. The `DataCompilation` process aligned the radiometer and housekeeping frames (Appendix C). I verified the alignment between receiver and mount data as part of the Routine Checks (§4.5.2).

The Time-code Reader provided the absolute time reference for the receiver¹⁸ by providing the decoded time (year, day, hours, minutes, seconds) to `adc_server`¹⁹. In addition, it created 1-Hz and 10-MHz clocks, phase-locked to the start of the GPS second, that synchronized the Master ADC. The 1-Hz clock had a 1/5 duty cycle, and the 10-MHz clock, a 1/2 duty cycle (square wave).

3.3 Control Software

The Control Software allowed the observers to implement simple commands and monitor the system status. The `Receiver Control Panel` provided high-level control of the receiver biasing. The `cbiviewer` process provided control and monitoring of the mount and commands for starting observations. Several stand-alone programs assisted in the creation of observing schedules. The `Live:QUIET` web page provided a graphical summary of the experiment status.

3.3.1 Receiver Control Panel

The `Receiver Control Panel` was a graphical interface for controlling and monitoring the receiver. It ran on a computer in the control room (control PC). It displayed the current

18. The mount electronics time-code reader received the same time code. Thus we synchronized the receiver and mount systems to each other and to the absolute GPS time.

19. Possibly useful information on configuring the Time-code Reader and its VME communication is in https://cmb.uchicago.edu/svn/kapner/quiet/crate_fs/branches/quadrature/root/vme_adc_threaded/src/TFP.c

state of the receiver, which could be

- **On** All biases were set for normal operation.
- **StandBy** The DAQ system was on, but the modules were off.
- **Off** The bias electronics and ADC Boards were off.
- **Down** All electronics, including the Receiver PC, were off.

We changed the state by clicking a button on `Receiver Control Panel`. It also provided a status display indicating whether the Receiver PC, ADC Backplane, SBC, WPID, and enclosure smart power strip were online; whether the bias electronics were on; and whether data were being properly recorded. [Chicago \(2010\)](#) provides further documentation.

3.3.2 Schedule Making

Several utilities assisted the observers in making the observing script each day. [Brizius \(2010\)](#) listed the astronomical sources to be observed. The `sched_maker` utility ([QUIET Collaboration, 2010](#)) provided a graphical display of when these sources would be within the elevation limits. The observers used it to generate a schedule file. The `QUIET Shift Utility` ([Nguyen, 2009](#)) checked that observations avoid the Sun²⁰. The `cbiviewer` process checked that the schedule file could be parsed. Ideally this check would find invalid schedules before they were submitted. However, logic errors (e.g. attempting to observe a source below the elevation limit) were not always detected in advance. Finally `cbiviewer` submitted the schedule to `cbicontrol`, which moved the mount to follow it.

20. We maintained at least 30° separation between the Sun and any source. The main worry was that radiation focused by the telescope could melt the cryostat vacuum window.

3.3.3 Real-time Monitoring

The web page `Live:QUIET`²¹ summarized the experiment status. It showed the last day of observation (mount motion and astronomical source), weather statistics, plots of the radiometer data from each module, receiver housekeeping, errors detected by the Online Software, and a picture from a digital camera pointing at the instrument (Figure 3.3). We performed a detailed check of this information three times per day (Tajima et al., 2010). We logged the findings in a Google spreadsheet²² in case we discovered a problem later.

3.4 Data Transfer and Storage

We transferred data regularly from the site to storage at the analysis centers. The process `cbicontrol` wrote data files²³ on the control PC in a proprietary format (“Level-0”). These files were automatically copied to the Storage PC in the control room, which had an attached RAID with space for 10 weeks of data. We copied newly created data files to Blu-ray disks (25 GB). Including downtime and compression, typical W-band observation generated 1–2 disks per day. We sent these disks to Chicago in weekly shipments. Upon arrival, we copied the data files to a file server with 29 TB of attached RAID storage. Furthermore, we copied the files by network to mirror servers in Oslo and the High Energy Accelerator Research Organization (KEK). Only when a data file existed in both Chicago and Oslo or KEK did we delete the file from the site. Thus we always stored two redundant copies of the data. Moreover, we used a Secure Hash Algorithm (SHA-1) checksum to confirm that the data was not corrupted in transit. (See also Tajima (2009c,a))

The Level-0 data files written by `cbicontrol` were not in a format optimized for analysis.

21. https://qufs.uchicago.edu/DQM/live_Wband/trunk/html/live2.html

22. [Online](#)

23. Typical files stored 15 minutes of data in 400 MB.



Figure 3.3: We took an image of the instrument every 5 minutes. The [Live:QUIET](#) web page displayed the latest image. We used the images for basic monitoring.

All data, whether receiver, mount, or housekeeping were interleaved because the major organization of the file was one frame per second of data. Moreover, different channels were grouped together by location (or multiplex address) in the readout electronics rather than by physical proximity. Conversely, a typical pattern in analysis was to access only a few channels that are in physical proximity for a long time interval. To make data files more amenable to such use, we automatically converted each Level-0 file to a corresponding Level-1 file.

We designed Level-1 data files for the typical analysis use case. The files used the common Flexible Image Transport System (FITS) format. Different parts of the experiment (receiver, mount, housekeeping, and biasing) were stored in separate Header Data Units (HDUs). Data from different diodes in the same module were stored together. Data from each channel were contiguous for time period of the data file (rather than 1s as in Level-0). (See also [Tajima \(2008\)](#))

3.5 Database

We stored metadata about each data file, scan, and observer log in a mySQL (Structured Query Language) database (Appendix [E.2](#)). For each data file we recorded its storage locations, begin and end time, and checksum. This information made it easy to check that no files were lost in the transfer process. For each scan we recorded the type (CMB, calibration, or other), source, status (successful completion or abort), and begin and end time. By combining the file and scan information we dynamically generated the list of files with data for any scan or set of scans. The daily observer log included the date, observer, and comments. Indexing the observer log made it search-able; this was useful when we discovered problems long after the observation.

Chapter 4

Observations

If I have not seen as far as others, it is because giants were standing on my shoulders.

Harold Abelson

We conducted W-band science observations between August 2009 and December 2010. QUIET deployed to the former CBI site in Chile’s Atacama Desert. During regular observation, we observed low-foreground regions of the sky every day and a weekly schedule of astronomical calibration sources. At the end of the season (December 2010) we performed special calibration observations with astronomical and artificial sources. Sometimes, noteworthy events interrupted the regular observation schedule. We routinely monitored the performance to correct problems as early as possible.

4.1 Site

The observation site was the Chajnantor Test Facility (CTF, $67^{\circ}45'42''$ W, $23^{\circ}1'41''$ S) on the Chajnantor Plateau in Chile’s Atacama Desert. CBI observed from the same site; we also used it for Q-band observations (QUIET Collaboration et al., 2011). The low atmospheric water vapor (Figure 4.1) and high elevation (5080 m, 50% of sea level pressure) reduced atmospheric emission and absorption: transmission was 98% in W band (QUIET Collaboration, 2012b; Pardo et al., 2001). Atmospheric fluctuations (“weather”) tended to be worse in the afternoon and the months December–March. The air temperature varied between -15°C (winter at night) and 10°C (summer days).

The site facilities supported observations. A heated, oxygen-enriched control room contained the Control and Storage PCs and stations for the observers. Two labs (one heated)

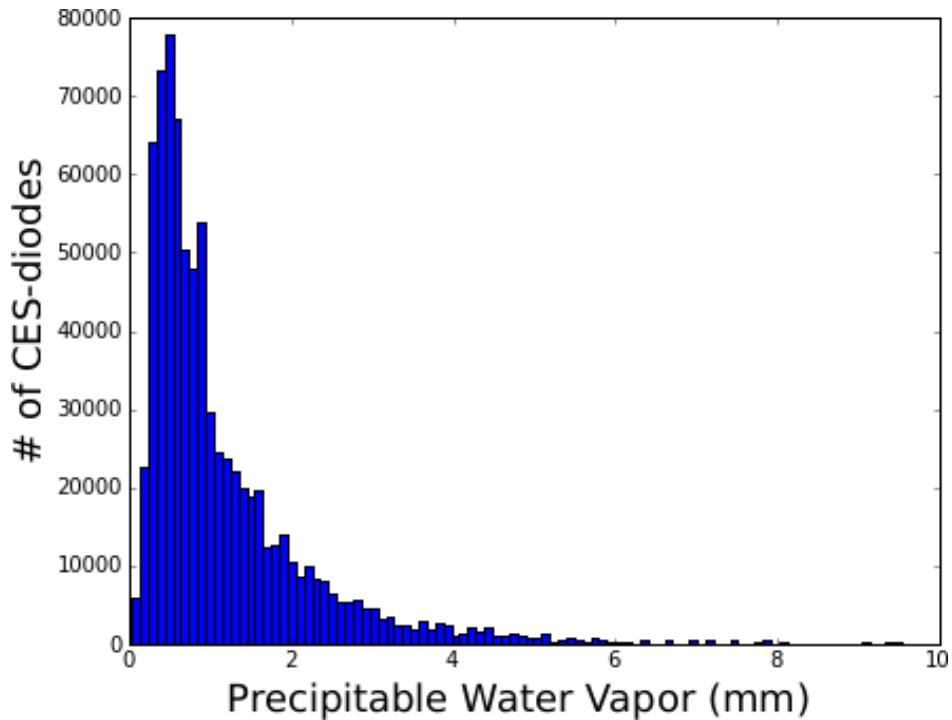


Figure 4.1: The precipitable water vapor (PWV) near the site (CTF) was typically < 1 mm. Low PWV corresponded to good weather. Our PWV monitoring came from the nearby Atacama Pathfinder EXperiment (Güsten et al., 2006). (I only show CES-diodes passing data selection.)

provided space for instrument investigation and repair. Two diesel generators provided power (60 kW) for the instrument. (One generator could power the entire site; the other was a backup.) Water chillers supported the water-cooled compressors needed to keep the cryostat cold. A fully retractable dome protected the instrument in bad weather (snow, extreme wind). Atacama Large Millimeter Array (ALMA) generously provided internet access. There were beds for the observers to sleep at the site; however, typically we returned to the nearby (1 hour) San Pedro de Atacama every night.

4.2 Regular Observations

We performed three types of regular observations: CMB, calibration, and Galactic. We observed the same four sky fields (“patches”) every day for CMB science. A weekly schedule of calibration observations supported the CMB observations by giving us a regular understanding of the instrument performance. Galactic observations are outside the scope of this thesis. We spent 71%, 15%, and 14% of the time on CMB, calibration, and Galactic observations, respectively (Table 4.1). The observing efficiency (observing time divided by calendar time) was 63%. We started regular observations on August 15, 2009 and finished on December 17, 2010.

4.2.1 CMB Observations

We regularly observed four patches of sky for CMB science (Figure 4.2 and Table 4.2). Each patch was roughly $15^\circ \times 15^\circ$. We picked the patches to be in low-foreground regions of the sky, determined from *WMAP* maps and dust templates (Samtleben, 2008). We also required the patches to be evenly distributed in Right Ascension (RA) so that at least one patch would always be above the elevation limit for observing (43°). Furthermore, no patch could be within 30° of the Sun. We observed patch 2a the most, and patch 7b, the least (Table 4.1). The total CMB observation time was 5336.9 hours. Patches 2a and 4a contained

Table 4.1. Observation Summary

Source	Hours Observed	Fraction of Observations (%)
CMB-1 (2a)	1855.0 ^a	25
CMB-2 (4a)	1443.6 ^a	19
CMB-3 (6a)	1388.7 ^a	19
CMB-4 (7b)	649.6 ^a	9
Total CMB	5336.9	71
Galaxy	1046.5 ^a	14
Moon	125 ^b	2
Tau A	324 ^b	4
Jupiter	168 ^b	2
Venus	5 ^b	0.07
RCW 38	31 ^b	0.4
Sky dip	250 ^b	3
Offset	190 ^c	3
Total Calibration	1090	15
Grand Total	7470	100

^aSOURCE: [Hasegawa \(2011b\)](#)

^bSOURCE: [Hasegawa \(2011c\)](#)

^cSOURCE: [Hasegawa \(2011d\)](#)

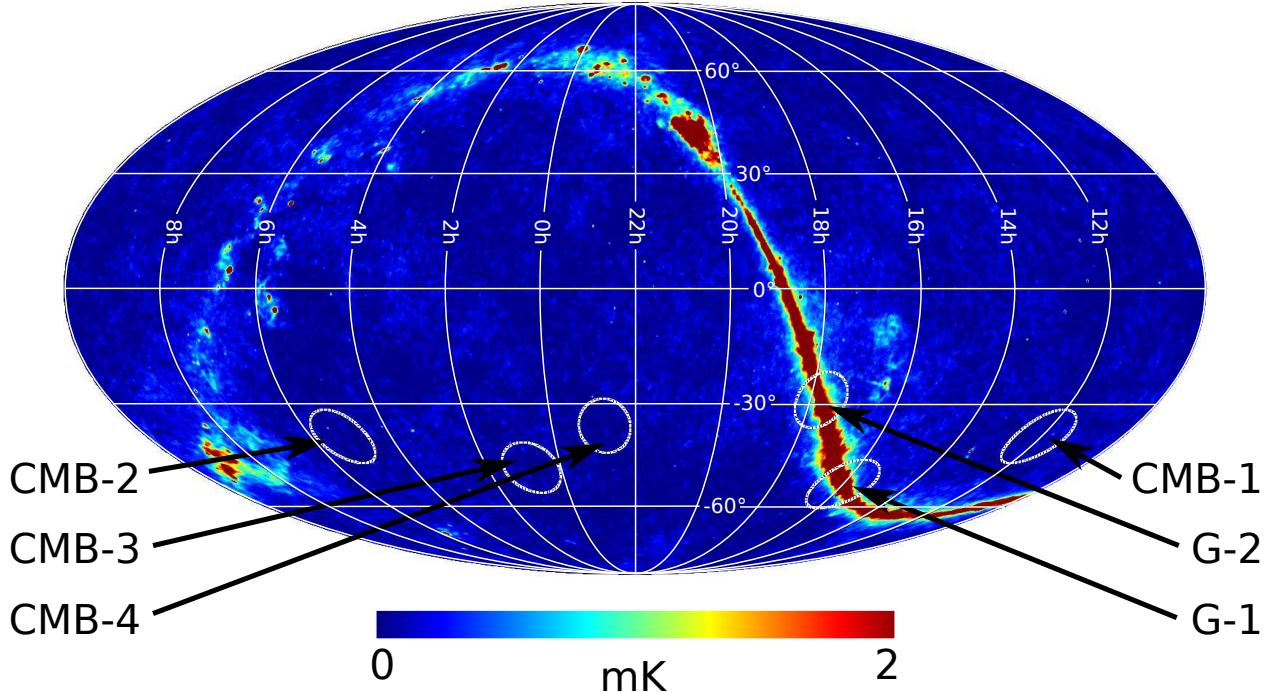


Figure 4.2: We observed four sky fields (“patches,” CMB-1–CMB-4) for CMB science. We chose the patches in low-foreground regions of sky. Two Galactic patches, G-1 and G-2, were for secondary foreground science goals. Coordinates are Right Ascension and Declination. The background is the *WMAP* Q-band temperature map.

significantly polarized astronomical sources: Centaurus A and Pictor A, respectively. We removed regions near these sources from analysis (§5.4).

When observing a patch, we conducted a series of constant-elevation azimuth scans (Figure 4.3). First, we pointed the telescope 7.5° (on the sky) ahead of the patch center in RA. Then we fixed the elevation and repeated azimuth scans. The scan amplitude was 15° (on the sky), and the period was $10\text{--}26\text{ s}$ ¹. The scan trajectory combined constant-jerk ($2.25^\circ/\text{s}^3$), constant-acceleration ($\leq 4.5^\circ/\text{s}^2$), and constant-velocity ($\leq 6.3^\circ/\text{s}$) segments (Figure 4.4). The typical scan speed was $5^\circ/\text{s}$ in azimuth. After ≈ 1 hour, the patch center drifted by 15° . At this point we ended the azimuth scans, pointed the telescope ahead of the patch again, and started a new series of azimuth scans. Each series of azimuth scans was a Constant

1. Limitations of the mount control required the scan period to be an integer multiple of 2 s (Shepherd, 2007, 2008b).

Table 4.2. CMB Patch Centers

Name	Right Ascension	Declination
CMB-1 (2a)	12 ^h 04 ^m	−39°00'
CMB-2 (4a)	05 ^h 12 ^m	−39°00'
CMB-3 (6a)	00 ^h 48 ^m	−48°00'
CMB-4 (7b)	22 ^h 44 ^m	−36°00'

Elevation Scan (CES). A uninterrupted (e.g. by the patch setting or a power failure) series of CESes of a single patch was a “run².²” Once the patch set below the elevation limit, we finished the run and began a new run of a different patch. Every week we rotated the deck angle by 45° so that observations would be uniformly distributed among deck angles.

4.2.2 Calibration Observations

We routinely observed astronomical sources to generate calibration data (Table 4.3). Many of the sources had limited availability so we interrupted CMB observations when necessary to observe the calibration sources.

We observed the supernova remnant Taurus A (“Tau A”) every day to calibrate the responsivity and beam shape. Each day we performed tight (< 1°) raster scans, centering two modules on Tau A. Every other day one of these modules was the central module, RQ45. The other module observations were evenly distributed across the array. We repeated each raster scan for five deck angles. Tau A observations took 75 minutes each day. Tau A was too close to the Sun between May 5 and August 6, creating a gap in observations (Nagai & Chinone, 2011b).

Every day we performed a large sky dip. The large sky dip was a periodic scan of the

2. There were a few exceptions. We defined the final CES boundaries after all observations were complete. In this process we discovered runs that were incorrectly assigned initially, and some parts of observations had to be discarded. As a result, we renumbered some runs in violation of the general rule.

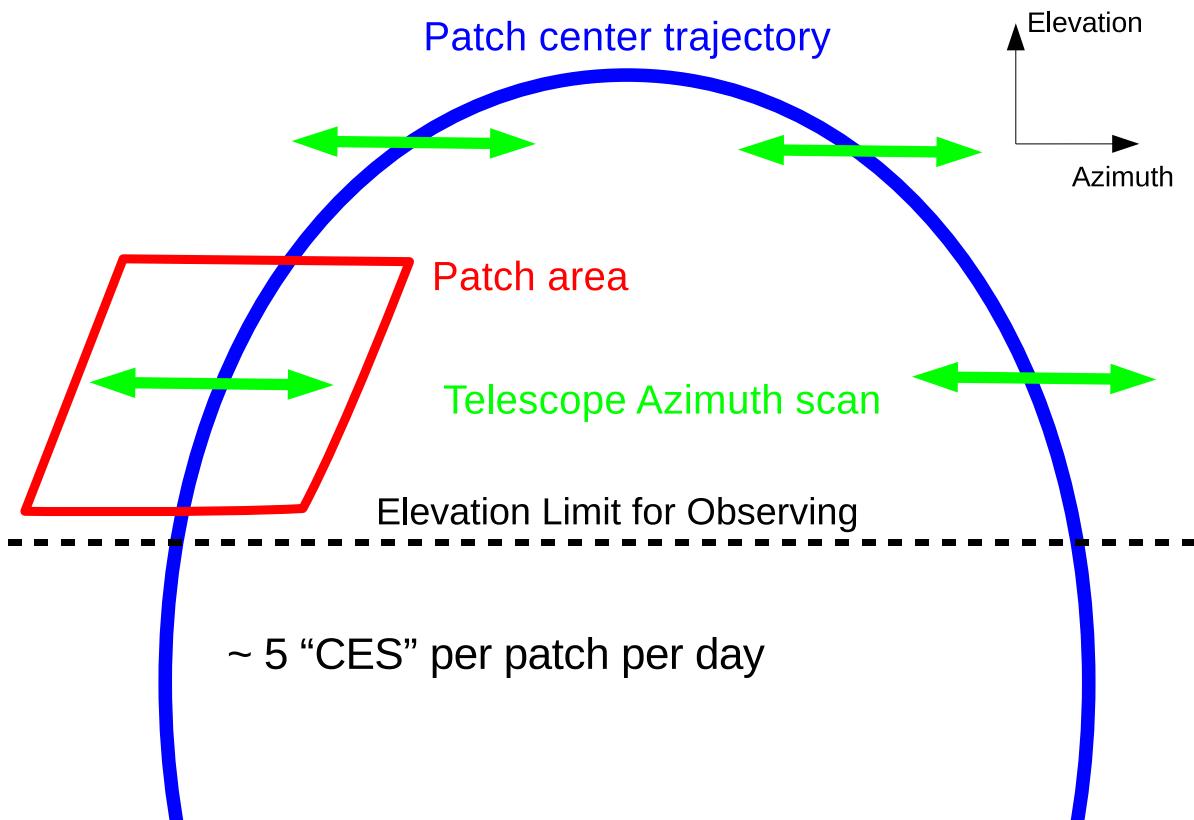


Figure 4.3: We observed each patch with a series of constant-elevation azimuth scans. We pointed the telescope ahead of the patch center in Right Ascension (RA) and fixed the elevation. The combination of a periodic azimuth scan and sky rotation mapped the patch area in two dimensions. After ≈ 1 hour (1 CES), we stopped the scanning, re-pointed the telescope, and repeated the process. In this way, we accumulated ≈ 5 CESes per patch per day.

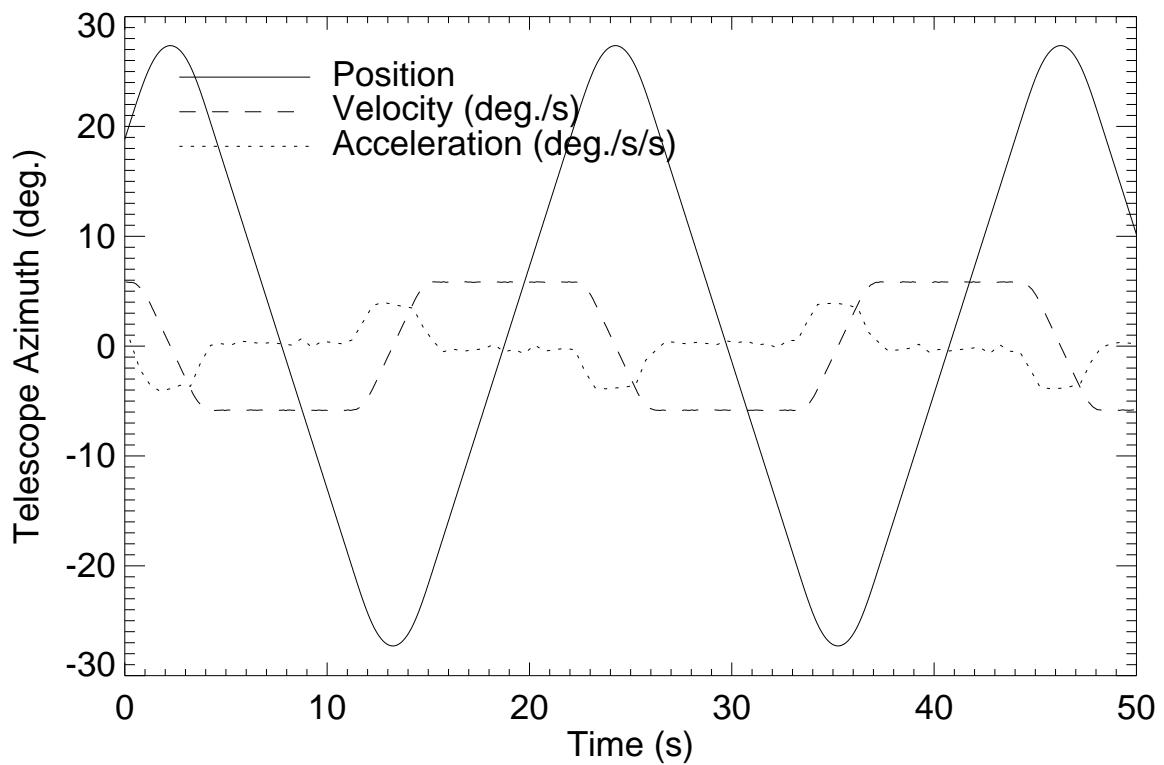


Figure 4.4: The CMB scan combined constant-velocity, constant-acceleration, and constant-jerk segments. This scan had period 22 s. Scans with sufficiently short period did not have constant-velocity segments.

Table 4.3. Regular Calibration Observations

Source	Observation Frequency	Duration (min.)	Calibration Obtained
Taurus A	1/day	75	Responsivity, beam shape, polarization angle
Large Sky Dip	1/day	15	Responsivity, $I \rightarrow Q/U$ leakage
Mini Sky Dip	1/CES	3	Responsivity, $I \rightarrow Q/U$ leakage
Moon	1/day	30	Pointing model
Moon	1/week	60	Pointing model
Jupiter	1/week	60	Pointing model, beam shape, responsivity
Jupiter	1/day	25	Pointing model, beam shape, responsivity
RCW 38	1/week	30	Pointing model, responsivity

telescope in elevation between 43° and 87° with azimuth and deck held fixed. The scan repeated six times and used 15 minutes. As the elevation changed, so did the effective atmospheric thickness and emission. We used this changing effective atmospheric temperature as the calibration source.

Before each CES we performed a mini sky dip and offset measurements. The mini sky dip, like the large sky dip, was a periodic elevation scan, but the amplitude was reduced to $\approx 5^\circ$. During offset measurements, we reverse-bias the phase switches to reduce their transmission. Then we measured the offset (voltage level with no RF power) of each diode. The DD offset did not influence the analysis because we only used the modulation of the polarization signal induced by scanning. However, the TP offset was necessary to relate TP to radiation intensity I .

We observed the Moon on both a daily and weekly schedule. Every day we performed a raster scan with the central row of the array³. The raster scan took 30 minutes. Every

3. RQ45 and whichever modules were aligned with it at the deck angle for CMB observations that week.

week we performed a full-array scan, with trajectory like that of a CES, that exposed every module to the Moon. We used the same full-array scan to observe Jupiter every week.

We used Jupiter, Venus, and RCW 38 (an HII region) to calibrate the differential-temperature modules. These observations contributed to the pointing model.

4.3 Season-end Calibrations

At the end of the season (December 2010) we performed special calibrations. We put a rotating sparse wire grid ([Tajima et al., 2012](#)) in front of the cryostat window. This calibrator created a modulated, polarized (1 K) signal in every module. We used the wire grid with varying cryostat and Electronics Enclosure temperatures to measure the dependence of the responsivity on the temperature. We used a swept, narrow-band source to measure the band-passes ([Chinone et al., 2012](#)). To map the far sidelobes, we performed several large angle scans using the Sun as a source. We performed large sky dips with additional window material in front of the cryostat to confirm the noise contribution from the window.

4.4 Noteworthy Events

Events throughout the season interrupted or otherwise affected normal observations. Table 3.3 of [Brizius \(2011\)](#) provides a complete list of such events. Here I list those that particularly impact this thesis:

1. On August 15, 2009 we replaced one of the Phase-switch Boards. This marked the beginning of regular W-band observing.
2. On 16 occasions we changed the CPID regulation parameters because the cryostat temperature was unstable.
3. On 6 occasions the generator failed, forcing a suspension of observations.

4. On 3 occasions we patched holes in the ground screen. We suspect these were related to residual ground pickup.
5. On 17 occasions the mount motors stalled in the middle of an observation, halting it.
6. On 11 occasions, degradation of the cryostat vacuum made it difficult to maintain the cryogenic temperature, and we suspended observations to pump down the cryostat.
7. On 8 occasions a power supply for one of the Preamp Boards failed. The result was unacceptable noise increases on a corresponding MAB. We replaced the power supply as soon as possible.
8. On 6 occasions we changed the Electronics Enclosure regulation parameters to keep the temperature stable.
9. Between September 26 and October 19, 2009, we replaced one of the cryostat cold heads because it exceeded the design lifespan and no longer provided sufficient cooling power.
10. On January 26, 2010 we installed the upper ground screen.
11. On May 30 and June 2, 2010, we found loose bolts in the cryostat window holder. We suspect these contributed to the vacuum leakage.
12. On December 17, 2010 we removed the upper ground screen for special calibrations. This marked the end of the W-band season.

4.5 Routine Checks

We monitored the instrument performance and data flow with routine checks every week. These checks supplemented the automatic real-time and daily checks (§3.3.3). The checks included data transfer and processing status, TP offset level, noise levels, responsivity, and

Table 4.4. Electronics Enclosure Temperature Regulation Setpoints

Time	Setpoint (°C)
August 15–November 29, 2009	35
November 29, 2009–April 18, 2010	40
April 18, 2010–October 30, 2010	35
October 30, 2010–December 17, 2010	40

polarization angle. I particularly checked the Electronics Enclosure temperature regulation and status of the Data Management system.

4.5.1 Enclosure Temperature

I monitored the Electronics Enclosure temperature throughout the season. Typically the temperature at the regulation point (in the Bias-Board Backplane) was stable to 1°C. However, the temperature was not uniform throughout the Backplane, and some bias boards (particularly those near the walls) experienced large (10°C) daily variations (Figure 4.5) following the ambient temperature. In summer, it became too hot during the day to keep the Enclosure regulated—the fan was on 100% of the time. At such times I increased the temperature setpoint. Table 4.4 lists the Electronics Enclosure setpoint temperatures during the season. With these adjustments, the typical temperature change for a bias board during a CES was < 1°C (Figure 4.6).

4.5.2 Data Management Status

I routinely checked the Data Management system to verify proper recording of the data. Possible problems included irregular timestamps, missing frames, de-synchronization of the Time-code Reader, Online Software crashes, power supply failures, loss of network connections, and ADC firmware errors. These problems could manifest incorrect alignment of

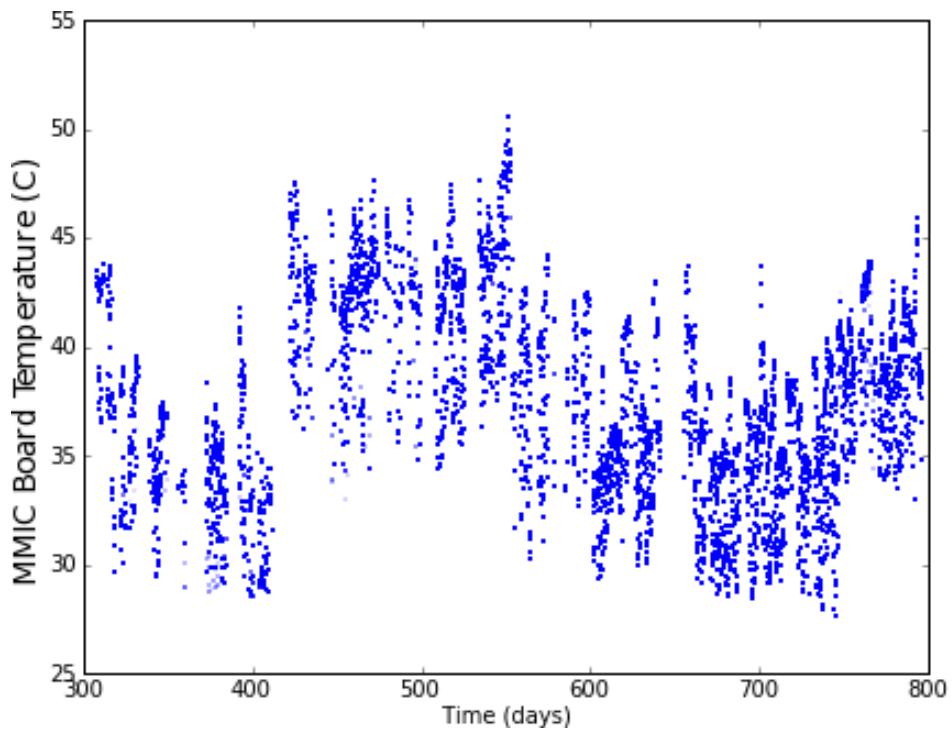


Figure 4.5: The Electronics Enclosure temperature changed by 20°C during the observing season. The change was due to two effects. Diurnal variation caused a 10°C oscillation each day. The regulation setpoint change throughout the season caused an additional 5°C change on month time scales. (I only show CESes that passed data selection.)

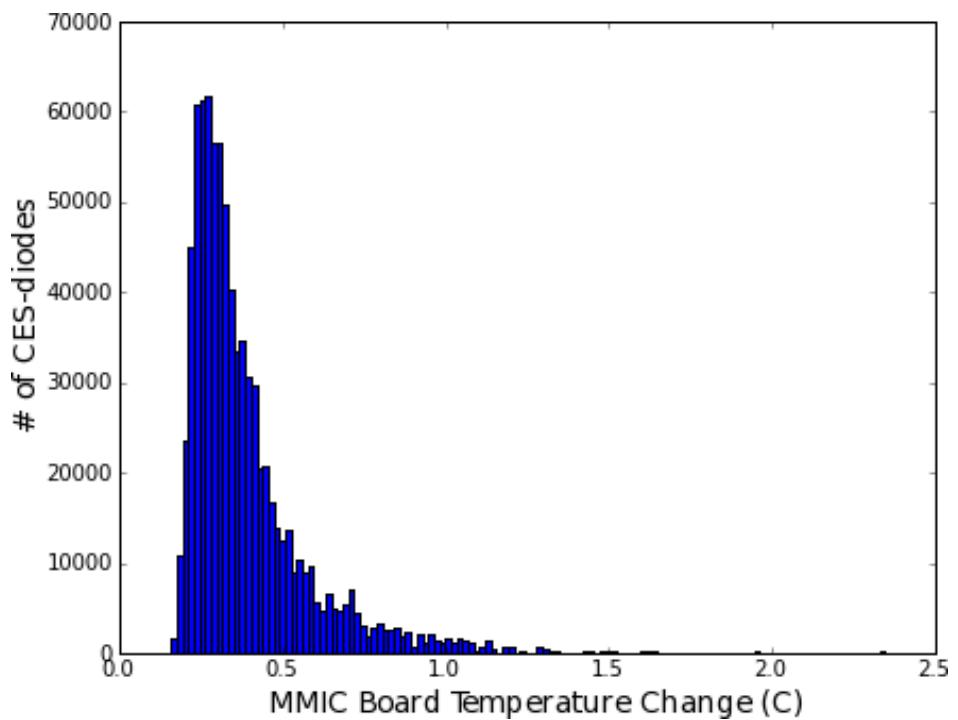


Figure 4.6: The Electronics Enclosure temperature was stable to $< 1^{\circ}\text{C}$ within a CES. The change for each CES is the standard deviation of all temperatures recorded within that CES. I only show CES-diodes that passed data selection.

receiver and pointing data, gaps in the data stream, and corrupt data values. I determined whether each CES was affected by these problems and whether or not we could use it (see w91_software in Appendix [E](#)). I also reviewed the status flags (Appendix [A](#)) for non-critical problems that nevertheless warranted action by the observers. The following critical problems caused CESes to be unusable:

1. Missing mount encoder data
2. Data length not an integer number of seconds
3. Contains a file with no data (this problem could be fixed by adjusting the CES definition)
4. Missing receiver status data
5. Missing receiver database entry
6. Missing receiver bias data
7. Missing receiver timestamps
8. Missing status flags
9. Missing housekeeping data
10. Missing receiver clock data
11. Uneven demodulation (the number of samples added and subtracted was not the same)
12. Incorrect number of accumulated 800-kHz samples in any 100-Hz sample
13. Frame counter mismatch between ADC Boards
14. Null frames (inserted to fix timing alignment problems, but invalid for properly aligned data)

15. BAD_TIMING flag (Appendix [A](#))
16. Missing frames (detected by frame counter not incrementing)
17. Timestamps not aligned to beginning of each second
18. Mismatch between radiometer, housekeeping, or mount timestamps
19. Timestamps not incrementing by expected interval
20. Irregularity in 50-Hz (double demodulation) clock
21. Irregularity in firmware 1-Hz clock
22. Incorrect blanking of phase-switch transitions
23. Incorrect housekeeping multiplexing

Chapter 5

Analysis

I might be blind, but I ain't stupid!

Ray Charles

Our analysis method¹ converted DD (§1.3.1) time streams into CMB polarization power spectra. To organize the analysis, we performed characterization and time-stream analysis separately on each CES, treating it as an independent analysis unit. The first characterization step was to calibrate the time-stream voltages into physical units. Then we modeled the noise in each CES to understand the statistical uncertainty. We filtered the data in the time domain to remove residual 1/f noise and ground contamination. We rejected data that had too much contamination or were otherwise unusable. After this processing in the time domain, we made maps and calculated the power spectra with a cross-correlation technique. To calibrate the analysis pipeline and understand its uncertainty, we performed the same analysis on hundreds of simulations of the entire experiment. We validated the analysis with a blind-analysis strategy. Following this strategy, we estimated systematic errors before knowing the power-spectrum results. We tried many analysis configurations (i.e. slightly different methods) before knowing the results. Unless otherwise noted, the following descriptions apply to the final configuration. I mention details of other, preliminary configurations when they had an impact on the final analysis choices. Appendix H lists all the configurations we tried.

The fundamental unit of analysis was the “CES-diode”: the data from one diode taken during one CES. For each CES-diode we calculated statistics for characterization of the data quality (§5.3.2 and Appendix E). These included noise, weather quality, housekeeping,

1. QUIET had two parallel analyses. This thesis describes only pipeline A (PCL). For discussion of the alternate analysis, see §6.2 and Naess (2012).

DAQ status, Type-B, and scan information. We also allowed the calibration parameters (e.g. responsivity) to vary on a CES-diode basis; although, only the pointing model and some Type-B parameters in fact changed between each CES. We performed filtering and data selection on a CES-diode basis; the smallest unit of data we accepted or discarded was a CES-diode. Only at the map-making stage did we combine CES-diodes.

5.1 Calibration

We needed to calibrate the instrument to understand the mapping between CMB polarization on the sky and time-ordered data (TOD) generated from the detectors. The relevant information was the beam shape, responsivity, detector angle, and a pointing model. An additional (compared to other CMB experiments) set of calibration parameters informed the algorithm we used to correct the Type-B differential non-linearity (§2.3). As much as possible we used astronomical calibration data to measure calibration constants, supplementing them with artificial calibrators when necessary. In general we considered two or more calibration methods and used the difference between them to calculate the calibration uncertainty.

5.1.1 Beam Shape

The beam shape (or point-spread function) effectively smooths the sky signal. We accounted for this smoothing in analysis (§5.4). We measured the beam shape from observations of Tau A and Jupiter as well as pre-deployment antenna-range measurements at Jet Propulsion Laboratory’s Mesa Antenna Measurement Facility². We used the Tau A measurements to set the beam FWHM and the antenna-range measurements to fix the shape beyond the central Gaussian peak. After radial averaging, the beam shape was approximately Gaussian with a significant non-Gaussian tail (Figure 5.1). We used a single, average beam shape

2. See [Monsalve \(2012a\)](#) for details on measuring the beam shape.

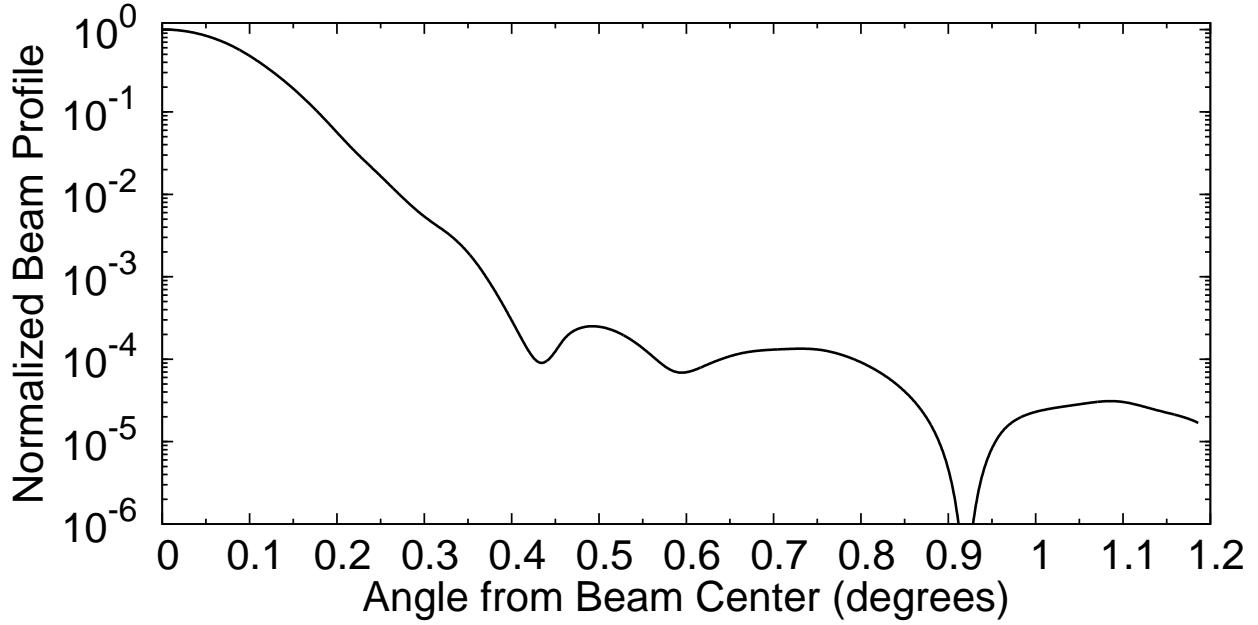


Figure 5.1: The beam shape was approximately Gaussian with a significant non-Gaussian tail.

for all modules; we did not detect any significant differences between the beams of different modules. We assigned systematic errors from the differences between Tau A measurements and optics simulations and from the differences between antenna-range measurements and Jupiter measurements with the differential-temperature modules.

The responsivity calibration sources were much smaller than the angular scales at which we measured the CMB. We corrected for this difference by scaling the calibration measurements to the angular scales of the CMB using the total beam solid angle, Ω . We calculated the solid angle by integrating the beam shape above, finding $\Omega = 13.63 \pm 0.65 \mu\text{sr}$. We propagated the systematic errors above. In addition, to estimate the error from the beam tail outside the measurement region, we compared the solid angle when integrating only the central $0^\circ.4$ instead of $1^\circ.4$ and assigned this difference as an additional (negligible) systematic error for the solid angle.

5.1.2 Responsivity

The responsivity was the coefficient of proportionality between the CMB fluctuations on the sky and the DD voltages. We modeled it as the product of two factors (Nagai, 2011a): the absolute responsivity was the coefficient for a single, well-characterized diode. The relative responsivity was the ratio between its responsivity and that of all other diodes.

The absolute responsivity factor in the model was the responsivity of the central module diode RQ45Q1. We measured it from repeated observations of Tau A. The maximum polarization signal was ΓP , where $P = 16.52 \pm 0.84$ Jy is the polarization flux (Nagai & Chinone, 2011a; Weiland et al., 2011) and

$$\Gamma = \frac{c^2}{2\nu^2\Omega k_B} \quad (5.1)$$

is the specific intensity to antenna temperature conversion factor. We calculated the responsivity from the ratio of this expected signal to the observed voltage. We found this responsivity had significant time dependence due to disruptions of the instrument. Therefore, we modeled the absolute responsivity as a piecewise-constant function of time. We fixed the time divisions based on jumps in the responsivity measurements that correlate with known disruptions of the instrument. Table 5.1 lists the time divisions. During the gap in Tau A observations (§4.2.2), we used large sky dips (§4.2.2) to track the time variation of the responsivity. We corrected the large-sky-dip data so the responsivities from large sky dips and Tau A were consistent when both were available (Nagai, 2011b). We then applied the antenna-to-thermodynamic-temperature conversion factor (1.25) so subsequent polarization measurements in K had thermodynamic (i.e. observing-frequency-independent) units³. The absolute responsivity uncertainty was dominated by uncertainty in Tau A's polarization flux (5.1%), extrapolation of RQ45Q1 responsivity to other diodes (3.8%, below), and uncertainty in the beam solid angle (4.8%). Uncertainty in the modeling of the time variation

3. Appendix E of Chinone (2011a)

Table 5.1. Absolute-responsivity–model Time Divisions

Start Time (UTC)	Responsivity (mV/K, thermodynamic)	Division Reason
	2.43	Beginning of observations
2009-10-19 09:14	2.64	Replaced cold head
2009-12-07 19:15	2.49	Fixed CPID oscillation
2010-03-02 21:37	2.39	Earthquake
2010-05-20 17:11	2.15 ^a	Cryostat warmed up
2010-05-29 22:39	2.23 ^a	Cryostat pumped
2010-07-28 22:54	2.27 ^a	CPID unstable
2010-08-02 12:22	2.27	Drain bias instability
2010-10-05 20:37	2.39	Cryostat pumped
2010-11-10 19:10	2.17	Generator failure

^aWe used large-sky-dip measurements for the responsivity.

was small (0.6%). We also considered dependence on the module physical temperature and electronics temperature, but it was negligible. The total absolute responsivity uncertainty was 8.3% (Nagai, 2012e).

We measured the relative responsivity from wire-grid and large-sky-dip observations. The relative responsivity of diode X to RQ45Q1 at time t is

$$R_X(t) = R_{X,\text{WG}} \frac{R_{X,\text{LSD}}(t)}{R_{X,\text{LSD}}(t_{\text{WG}})}, \quad (5.2)$$

where t_{WG} is the time of the wire-grid measurement, $R_{X,\text{WG}}$ is the responsivity ratio between diode X and RQ45Q1 as measured from the wire grid, and $R_{X,\text{LSD}}(t)$ is the responsivity ratio measured from large sky dips. We used the median of large-sky-dip measurements in each time division when constructing relative responsivities (Ishidoshiro, 2011). To estimate the uncertainty in the relative responsivity, we compared our model to the Tau A measurements

for off-center modules⁴. The mean responsivity shift from this comparison was 1.2%⁵, and the RMS scatter was 12.5%. If we used only Tau A measurements to model the responsivity, the result would change by 3.8%.

5.1.3 Detector Angle

The detector angle was the polarization-sensitive axis of each diode. Typically⁶ the Q1 and Q2 diodes had 90° between their detector angles. The U1 and U2 diodes were separated by 90°, and each Q diode was offset by 45° from each U diode. However, the overall angle remained free⁷, and imperfections might alter the relationship between the diodes. Therefore, we measured the axis for each diode. As with responsivity, we separated the calibration of detector angle into absolute and relative components.

We measured the absolute detector angles for RQ45 from the known polarization direction of Tau A (Aumont, J. et al., 2010; Chinone et al., 2011). Each day observation yielded a precision of 3°. By combining all measurements, the total measurement uncertainty was 0°.3. The known polarization direction itself had an uncertainty of 0°.2. The possible error in extrapolating the angle to non-central modules (see below) was 0°.2. We also considered possible variation of the detector angle with applied bias, and limited the effect to 0°.2. Combining these uncertainties, the total uncertainty on the absolute detector angle was 0°.5.

We measured the angles of all non-central detectors relative to RQ45Q1 from the wire-grid season-end calibration. This calibration simultaneously illuminated all detectors with a modulated polarization signal. From the relative phases of different diodes, we calculated the relative angles between their axes (Tajima, 2010). We compared the resulting angles to those obtained from the small number of off-center Tau A measurements. The resulting

4. Except for RQ45Q1, the model does not depend on Tau A measurements.

5. §3.1 of Nagai (2012e)

6. i.e. if there was no hybrid-coupler imperfection

7. It was affected by extra phase lengths in the module.

mean shift was $0^\circ.2$, and we included it in the absolute angle uncertainty above. The root-mean-square (RMS) scatter was $3^\circ.1$. The uncertainty of the wire-grid measurement was $0^\circ.8$. Combining them, we assigned a total relative uncertainty (i.e. scatter among the diodes) on the polarization angle of $3^\circ.2$.

5.1.4 Pointing

The pointing model converted the mount encoder positions into angular positions on the sky for each module. The model had two components: a mount model specifying the position of the central module and an array model specifying the relative positions of the other modules.

We constructed the mount model from observations of Jupiter, RCW 38, and the Galactic center with the differential-temperature modules ([Oslo, 2012](#)). For each observation we had the apparent position (encoder measurements) and true position (ephemeris). The mount model mapped between them, and we chose the parameters of the model to minimize the difference between the true positions and the positions as reconstructed by the model. The model itself was a series of rotation matrices,

$$M_{\text{mount}} = M_{\text{eltilt}} M_{\text{aztilt}} M_{\text{bore}} M_{\text{coll}}. \quad (5.3)$$

Using the zyz Euler convention and expressing each matrix as a series of rotation angles

$M(\alpha, \beta, \gamma)$, the matrices are

$$M_{\text{eltilt}} = M(-(90^\circ + \text{az}), \theta_E, 90^\circ + \text{az}) \quad (5.4)$$

$$M_{\text{aztilt}} = M(-\psi, \theta, \psi) \quad (5.5)$$

$$\begin{aligned} M_{\text{bore}} = & M(-\text{az} - (\Delta_{\text{az}} + \theta_{\text{ef,az}} \cos(-\text{az} + \psi_{\text{ef}})), \\ & 90^\circ - \text{el} - (\Delta_{\text{el}} + \theta_{\text{ef,el}} \cos(-\text{az} + \psi_{\text{ef}})), \\ & \text{dk} - (\Delta_{\text{dk}} + \theta_{\text{ef,dk}} \cos(-\text{az} + \psi_{\text{ef}}))) \end{aligned} \quad (5.6)$$

$$M_{\text{coll}} = M(\psi_c + 180^\circ, \theta_c(1 + \theta_{\text{ec}} \cos(\text{dk} + \psi_{\text{ec}})), -\psi_c - 180^\circ), \quad (5.7)$$

where $(\text{az}, \text{el}, \text{dk})$ are the azimuth, elevation, and deck encoder positions. Table 5.2 gives the values and physical meaning of the model parameters.

We performed several tests to evaluate uncertainty in the mount model:

1. We combined all differential-temperature data to make a map of the source PMN J0538-4405 in CMB patch 4a. The difference between the true (Gold et al., 2011) and reconstructed positions of this source was $0.4 \pm 1'$. The FWHM beam size measured was $14.2' \pm 1'$. The increase compared to the beam size found above ($11.66'$) was due to scatter in the pointing model.
2. We measured this scatter from the residuals in the pointing model fit. We fit a two-dimensional (azimuth and elevation) Gaussian to the residuals, and used the difference between the major and minor axes to estimate the uncertainty. We found that the pointing model had an intrinsic scatter FWHM of $5.1 \pm 0.25'$. We included this scatter as an additional factor in the beam deconvolution (§5.4).
3. As a check of this additional effective beam, we used the differential-temperature data to calculate the CMB temperature power spectrum, and compared the result to a Λ CDM model. This comparison confirmed the additional beam smoothing with an

Table 5.2. Pointing-model Parameters

Name	Meaning ^a	Value
Δ_{az}	Azimuth encoder offset	$-0^\circ 03547$
Δ_{el}	Elevation encoder offset	$-0^\circ 02938$
Δ_{dk}	Deck encoder offset	$0^\circ 10832$
ψ	Azimuth tilt orientation	$49^\circ 62021$
θ	Azimuth tilt magnitude	$-0^\circ 00701$
θ_E	Elevation tilt	$-0^\circ 00371$
ψ_C	Collimation offset direction	$26^\circ 67075$
θ_C	Collimation offset magnitude	$0^\circ 29481$
ψ_{ec}	Collimation ellipticity orientation	$253^\circ 82681$
θ_{ec}	Collimation ellipticity magnitude	0.2126^b
ψ_{ef}	Encoder offset flexure direction	$11^\circ 35199$
$\theta_{\text{ef,az}}$	Encoder offset flexure magnitude in azimuth	$-0^\circ 02252$
$\theta_{\text{ef,el}}$	Encoder offset flexure magnitude in elevation	$-0^\circ 00424$
$\theta_{\text{ef,dk}}$	Encoder offset flexure magnitude in deck	$-0^\circ 26147$

^aSee also [Shepherd \(2008c\)](#); [Bischoff \(2009\)](#).

^bThis parameter is not an angle.

uncertainty of ± 1.27 .

4. We also used Moon observations to test the pointing model. The RMS scatter between the position predicted by the model and the actual position was $1.6'$ ([Kusaka, 2012a](#)).

We constructed the array model from Moon observations. In a full-array Moon observation, all modules observed the Moon. From the time delay between the Moon signal in different modules, we calculated the relative position of each module ([Figure 1.4](#) and [Kusaka \(2011b\)](#)). The uncertainty as measured by the scatter between different Moon observations was $< 0.2'$. We also constrained the offset of the deck angle to $< 0.1^\circ$.

8. During part of Q-band observation the deck encoder was loose (§4.2 of [QUIET Collaboration et al. \(2011\)](#)). Therefore, we wanted to confirm this problem did not occur in W band.

5.1.5 Type-B-glitch Correction

We corrected the data affected by Type-B glitching. We based the correction algorithm on a phenomenological understanding of the glitches. The algorithm had several free parameters to fit or measure. I address the uncertainty introduced by this correction in §5.7.3.

We based the correction on a Type-B glitch model (Kusaka, 2010c). Suppose that the voltage at the ADC input was x^i . Then we modeled the ADC operation as

$$y^i \equiv f(x^i) \equiv x^i + \Delta f(x^i), \quad (5.8)$$

where

$$\Delta f(x) = \sum_i \Delta f_i(x) \quad (5.9)$$

encodes the effect of glitches, and each glitch is

$$\Delta f_i(x) = \begin{cases} -A_i/2 & \text{if } x < x_0^i \\ 0 & \text{if } x = x_0^i \\ A_i/2 & \text{if } x > x_0^i, \end{cases} \quad (5.10)$$

where x_0^i and A_i are the location and height of the i th glitch. The correct DE and TP data are

$$x_{1(2)} = \frac{1}{N} \sum_{i=1}^N x_{1(2)}^i \quad (5.11)$$

$$x_{\text{DE}} = \frac{x_1 - x_2}{2} \quad (5.12)$$

$$x_{\text{TP}} = \frac{x_1 + x_2}{2}, \quad (5.13)$$

where subscripts 1 and 2 indicate samples in the two 4-kHz phase-switch states, and $N =$

3440 is the number of 800-kHz samples per 100-Hz sample per phase. However, the ADC output is

$$y_{1(2)} = \frac{1}{N} \sum_{i=1}^N y_{1(2)}^i \quad (5.14)$$

$$y_{\text{DE}} = \frac{y_1 - y_2}{2} \quad (5.15)$$

$$y_{\text{TP}} = \frac{y_1 + y_2}{2}. \quad (5.16)$$

Since the 800-kHz noise was Gaussian (§5.7.3), we define

$$y = F(x; \sigma) \equiv \int f(s) P_G(s; x, \sigma) ds = x + \int \Delta f(s) P_G(s; x, \sigma) ds, \quad (5.17)$$

where $P_G(s; \mu, \sigma)$ is a Gaussian distribution

$$P_G(s; \mu, \sigma) \equiv \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(s - \mu)^2}{2\sigma^2}\right) \quad (5.18)$$

centered at μ with standard deviation σ . The inverse function

$$G(y; \sigma) \equiv F^{-1}(y; \sigma) \quad (5.19)$$

satisfies

$$x = G(F(x; \sigma); \sigma). \quad (5.20)$$

We estimate corrected values

$$x'_{1(2)} = G(y_{1(2)}; \sigma). \quad (5.21)$$

In case the correction parameters are accurate and $N \rightarrow \infty$, $x' = x$. I discuss departures from these idealities later. Defining $\Delta G(y; \sigma) \equiv G(y; \sigma) - y$, we can calculate the corrected

data x' if we know the function ΔG . From Eq. 5.17,

$$\begin{aligned}\Delta F(x; \sigma) &\equiv F(x; \sigma) - x = \\ \int \Delta f(s) P_G(s; x, \sigma) ds &= \sum_i A_i \int_0^{x-x_0^i} P_G(s; 0, \sigma) ds = \sum_i \frac{A_i}{2} \operatorname{erf} \left(\frac{x-x_0^i}{\sqrt{2}\sigma} \right).\end{aligned}\quad (5.22)$$

Because

$$x = G(y; \sigma) \quad (5.23)$$

$$= y + \Delta G(y; \sigma) \quad (5.24)$$

$$= F(x; \sigma) + \Delta G(y; \sigma) \quad (5.25)$$

$$= x + \Delta F(x; \sigma) + \Delta G(y; \sigma), \quad (5.26)$$

$$\Delta F(x; \sigma) = -\Delta G(y; \sigma). \quad (5.27)$$

Using Eq. 5.27 we solve for ΔG using Newton's method. If

$$h(s) \equiv s + \Delta F(s + y; \sigma), \quad (5.28)$$

then the value of s that solves $h(s) = 0$ is ΔG . Thus we solved $h(s) = 0$ with Newton's method to calculate ΔG and applied it to y to calculate the corrected data x' . (An additional subtlety is that the measurement of glitch locations x_0^i included glitching. See Appendix A.3.1 of [Kusaka \(2010c\)](#) for a solution.)

The algorithm above had free parameters x_0^i (glitch locations), A_i (glitch heights), and σ (noise level). We measured the glitch locations and heights by injecting known voltage signals into each ADC channel ([Brizius, 2009](#)). Although there were several patterns, not all channels had the same glitch locations and heights. Some channels did not have any glitching. Because the measurement did not cover the full ADC range, I extrapolated the patterns to

cover the range used in observations (Buder, 2011g). Moreover, I modified the parameters for a small number (5) of channels for which the ADC measurement was inconsistent with glitching found in CMB data. Because σ is the 800-kHz noise level, the relationship between it and the 100-Hz noise depends on the Preamp Board filtering. A calculation from the Preamp Board circuit design gave 0.6342 as the ratio between the noise levels. However, using that to calculate σ did not result in good correction. Instead, I estimated σ from CMB data in an iterative procedure.

As an intermediate step in calculating σ , I defined a “preamp factor” for each diode. This was the ratio between 800-kHz and 100-Hz noise levels that I estimated empirically. For each CES-diode I calculated a “Type-B χ^2 ” (§5.3.2) that quantified the amount of Type-B glitching in it. First I selected CESes with bad Type-B (ordered by the difference in Type-B χ^2 after correction), but otherwise good data quality. For each diode I selected the worst 20 such CESes. For each such CES-diode, I found the preamp factor that minimized the Type-B χ^2 . Then I averaged the preamp factors that made the $\chi^2 < 3$ and came from CES-diodes with significant Type-B glitching⁹. I calculated an uncertainty for the preamp factor from the scatter among CESes. For a small number (6) of diodes, I could not find any CES from which I could estimate a preamp factor. For such diodes, I used the average preamp factor of the well-measured diodes.

Once I fixed the preamp factor, I calculated σ for each CES-diode in an iterative procedure. First I made a preliminary guess for the noise level from “pseudo-quadrature” 25-Hz data, in which I differenced adjacent DD samples. These pseudo-quadrature data had the same noise level but no signal. The noise level guess was the standard deviation of the pseudo-quadrature data. Then I corrected the CMB data using the guess for σ . From the corrected data, I computed a new noise level guess and used it to correct the data again. I iterated this process until the noise level converged to $0.0002 \text{ bit}\sqrt{s}$. (Typical diode noise

9. The glitching was significant if $\Delta\chi^2 > 1$. Some diodes rarely manifested Type-B; for such diodes the worst 20 CESes may contain CESes that did not manifest Type-B.

levels were $0.2 \text{ bit}\sqrt{\text{s.}}$)

Because we did not record 800-kHz data all the time, the correction had a statistical uncertainty

$$\delta x_{\text{DE}} = x'_{\text{DE}} - x_{\text{DE}} = \frac{\delta x_1 - \delta x_2}{2}, \quad (5.29)$$

where

$$\delta x \equiv G(y; \sigma) - x \quad (5.30)$$

is the miscorrection, in this case due to statistical fluctuation.

$$\sqrt{<\delta x_{\text{DE}}^2>} \lesssim \frac{A}{2\sqrt{N}}. \quad (5.31)$$

For the largest glitches with $A \approx 40$, this noise was $0.36 \mu\text{V}\sqrt{\text{s}}^{10}$. The typical noise level was $3 \mu\text{V}\sqrt{\text{s}}$ so this Type-B correction noise was not dominant, and we included it when we measured the noise from the data after correction. I evaluate the systematic errors due to Type-B in §5.7.3.

5.2 Noise Modeling

We modeled the noise in each CES-diode as a combination of white and 1/f noise components. We fit the model

$$\hat{N} = \sigma^2 (1 + (f_0/f)^\alpha) \quad (5.32)$$

to each CES-diode DD data in the Fourier domain. The model parameters were σ , the white noise level, f_0 , the 1/f knee frequency, and α , the 1/f slope. The fitting excluded frequencies between 0.8 and 1.2 times the scan frequency as well as frequencies above 9.6 Hz (the low-pass filter cutoff, §5.3.1). We included correlations among the diodes of each module (Buder & Kusaka, 2011). To study these correlations, we computed the covariance matrix at each

10. Appendix A.4.1 of Kusaka (2010c)

frequency

$$\hat{C}_{ij}(f) = F_i(f)F_j^*(f), \quad (5.33)$$

where $F_i(f)$ is the Fourier transform of diode i . At high frequencies the white noise dominated so

$$\hat{\rho}_{ij}^w = \Re \left[\left\langle \frac{\hat{C}_{ij}(f)}{\sqrt{P_i(f)P_j(f)}} \right\rangle_{2-4.5 \text{ Hz}} \right] \quad (5.34)$$

estimates the white noise correlation matrix with minimal bias from residual 1/f noise when we average between 2 and 4.5 Hz. P_i is the noise power ($F_iF_i^*$). We took the real part after checking that the imaginary part of the correlation matrix was negligible. I also checked that inter-module correlations were negligible. In two CESes the largest inter-module correlation coefficient was 0.25, and only one pair of modules had such a large correlation. The expected correlation matrix for an ideal module is

$$\rho_{ij}^w = \begin{pmatrix} 1 & 0.5 & 0.5 & 0 \\ 0.5 & 1 & 0 & 0.5 \\ 0.5 & 0 & 1 & 0.5 \\ 0 & 0.5 & 0.5 & 1 \end{pmatrix}, \quad (5.35)$$

where the order of the rows and columns is (Q1, U1, U2, Q2). However, module imperfections altered the off-diagonal elements ([Bischoff & Kusaka, 2009](#)). The median measured correlation matrix was ([Kusaka et al., 2012](#))

$$\hat{\rho}_{ij}^w = \begin{pmatrix} 1 & 0.495 & 0.502 & 0.058 \\ 0.495 & 1 & 0.061 & 0.516 \\ 0.502 & 0.061 & 1 & 0.387 \\ 0.058 & 0.516 & 0.387 & 1 \end{pmatrix}. \quad (5.36)$$

Note that correlation between Q and U diodes did not cause significant correlation between measurements of Q and U because we effectively differenced $Q1$ from $Q2$ in map making. The difference $Q1 - Q2$ is uncorrelated with $U1 - U2$ for any correlation matrix.

5.2.1 Correlated 1/f Noise Investigation

We investigated the 1/f noise correlations in several CESes before deciding on the treatment of them in the noise model. This section details our investigation and findings; §5.2.2 describes the final noise model. First we defined a 1/f correlation estimator for a single CES. Then I diagonalized the 1/f correlation matrix for each CES. I found some significant correlated modes and confirmed the significance with simulations. Then we showed that for frequencies above the pipeline filter cutoff, the correlated modes do not contribute significantly to the total noise. We concluded that we could ignore these correlated modes in the noise model, thereby making at most a 1% error in noise estimation.

Measuring the 1/f noise correlation was more complicated because even at low frequencies the white noise may not be negligible. Moreover there can be multiple 1/f components with different frequency dependencies. We used

$$\hat{\rho}_{ij}^{1/f} = \left(\sum_f \frac{(C_{ij}(f) - \langle C_{ij}(f) \rangle_{2-4.5\text{Hz}}) \sqrt{w_i w_j}}{\sqrt{\sigma_i^2 \sigma_j^2 \left(\frac{f_0}{f}\right)^{\alpha_i + \alpha_j}}} \right) / \sum_f \sqrt{w_i w_j} \quad (5.37)$$

as a 1/f correlation matrix estimator for each CES. The w_i are weights chosen to optimize the signal to noise

$$w_i(f) = \sigma_i^2 \left(\frac{f_0}{f}\right)^{\alpha_i} / \hat{P}_i(f). \quad (5.38)$$

We excluded frequencies between 0.8 and 1.2 times the scan frequency (because there was narrow-band noise from the scan) and above 1.5 times the scan frequency.

To study the correlation structure, I diagonalized $\hat{\rho}_{ij}^{1/f}$ for several CESes. Before actually

performing the matrix diagonalization, I took several steps to regulate the “raw” correlation matrix $\rho_{ij}^{1/f}$. First, I removed several diodes. All known non-functional diodes were cut. Diodes with knee frequency $< 10\text{ mHz}$ were removed because there was very little $1/f$ noise so the correlation coefficient had large fluctuations. I also cut diodes with small (< 0.1) diagonal elements in the correlation matrix. Since I used a fit to estimate the variance, it is not guaranteed that the diagonal elements of the raw correlation matrix are 1. Although I renormalized the matrix (below) so this is the case, extremely small diagonal values here could cause the normalization to become unstable. The normalized correlation matrix is

$$\rho_{ij}^{\text{normalized}} = \frac{\hat{\rho}_{ij}^{1/f}}{\sqrt{\hat{\rho}_{ii}^{1/f} \hat{\rho}_{jj}^{1/f}}} \quad (5.39)$$

After normalization, I forced off-diagonal elements to be in the range $[-0.99, 0.99]$ by reducing the absolute value of outliers to 0.99. The normalization process ensured that the diagonal elements were 1. After diagonalization, I obtained the eigenvalues and eigenvectors. Small negative eigenvalues remained even after regulation (Figure 5.2). These were an effect of the difference between the true correlation matrix (which must be positive definite) and the estimated correlation matrix. For typical CESes, the first mode had a significant coefficient for most diodes. That is, the mode with the largest $1/f$ variance was correlated among most diodes in the array. Figure 5.3 shows the first mode coefficient vs. $1/f$ slope. The correlation increased with slope, and all diodes with slope $\gtrsim 1.5$ were correlated. This was evidence for a two-component model of $1/f$ noise with the slope-2 component correlated among all diodes. We attributed the slope-2 component to atmospheric fluctuations and the slope-1 component to MMIC-amplifier-gain fluctuations.

Since there was statistical noise in the correlation measurement (due to finite data size), I ran a toy simulation of one CES to ensure the measurement was significant. In the simulation I assumed there were no correlations and that the noise was Gaussian. A typical CES

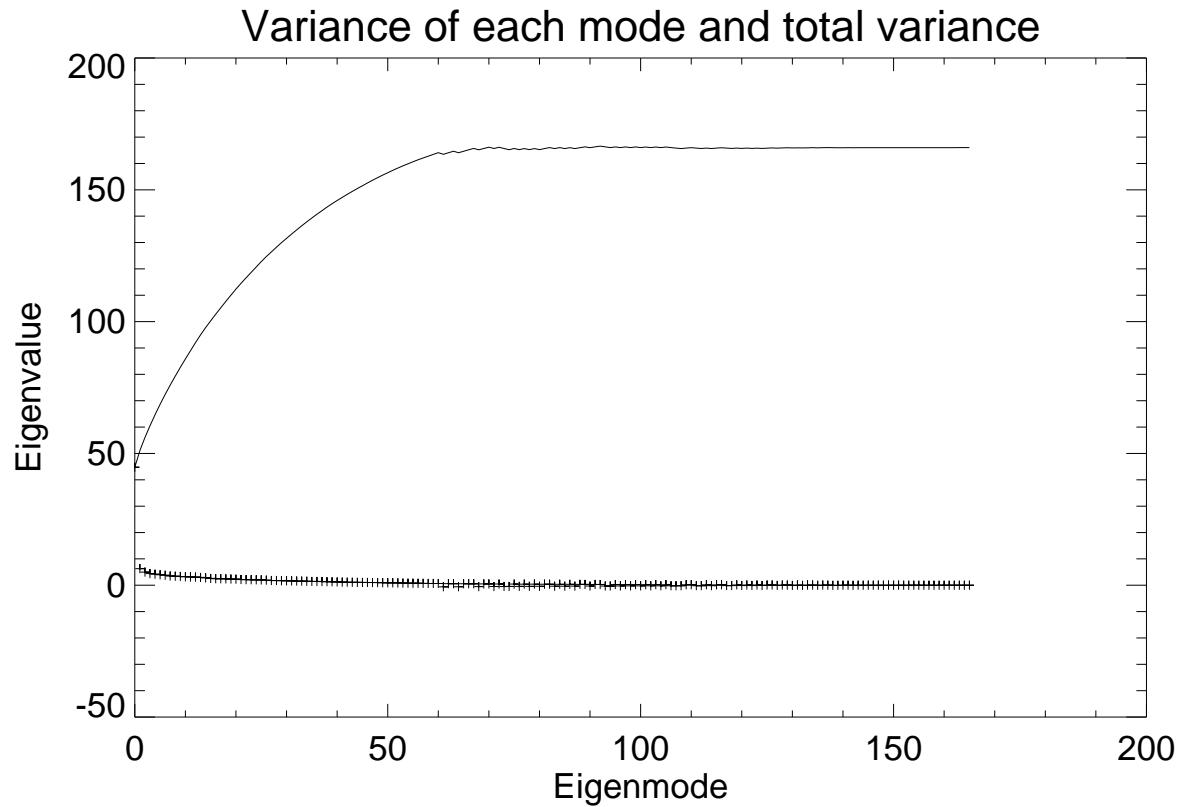


Figure 5.2: A small number of modes contained most of the $1/f$ variance. The crosses are the eigenvalues of the $1/f$ correlation matrix for one CES. I ordered them by absolute value. The solid line is the sum of the first N eigenvalues. The negative eigenvalues (crosses slightly below 0 and points where the solid line decreased) were due to fluctuations; with an infinite data length the correlation matrix would be measured perfectly, and there would be no negative eigenvalues.

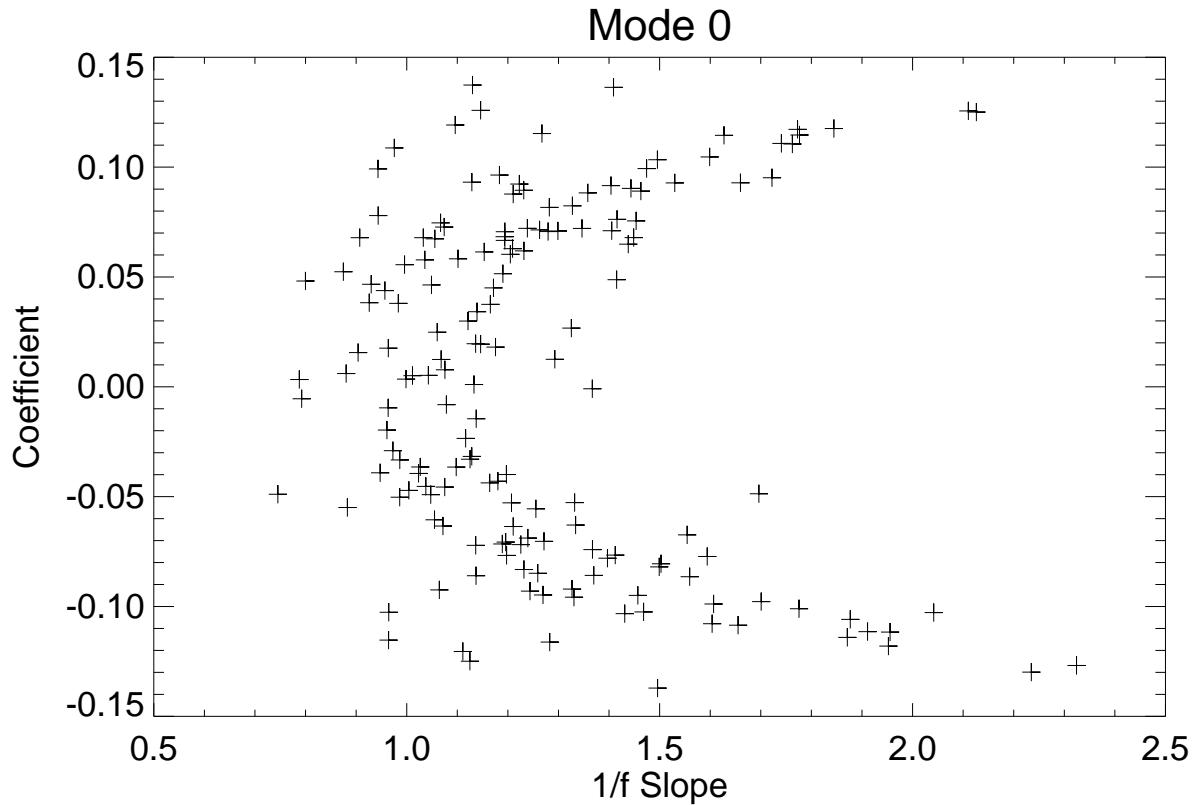


Figure 5.3: The coefficients of the first eigenvector of the $1/f$ correlation matrix were correlated with the $1/f$ slopes of each diode's noise spectrum. Each point shows the $1/f$ slope and eigenvector coefficient for one diode; all points are from the same CES. There were two $1/f$ components, and the slope-2 component, which we attributed to weather, was correlated among all diodes. The slope-2 component generated the first eigenmode. When that component dominated the $1/f$ noise of a diode, that diode was strongly correlated with the rest of the array. Thus the diodes with measured slope ≈ 2 had consistently high coefficients in the first eigenvector. We attributed the slope-1 component to MMIC-amplifier-gain fluctuations.

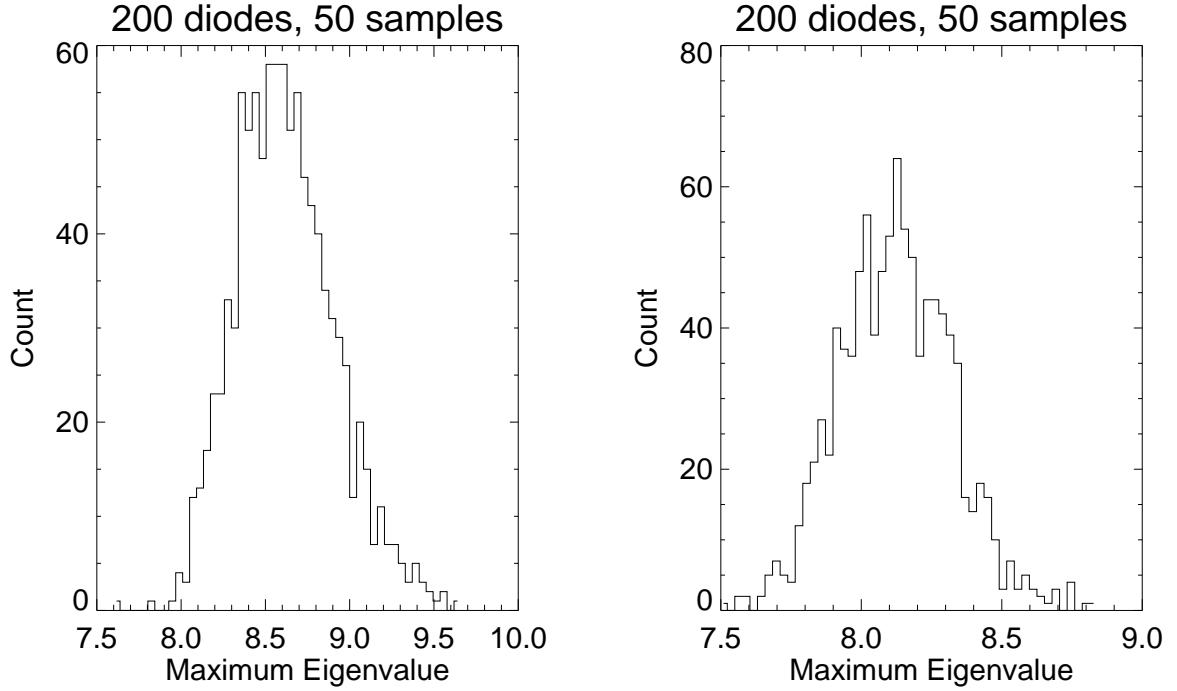


Figure 5.4: I confirmed that the $1/f$ correlation measurement was significant. The left (right) panel shows the first (second) eigenvalues from simulations. When the $1/f$ correlation was significant, the data eigenvalues were much larger (≈ 50).

was 5000 s so there were 5000 Fourier modes / Hz. I only considered diodes with a knee frequency above 10 mHz, and the weight w was small above the knee frequency. Therefore, I used 10 mHz as an approximation to the bandwidth used in the correlation measurement. I simulated 50 independent Gaussian random samples for each diode and computed the resulting correlation matrix. I diagonalized the simulated correlation matrix and computed the distributions of the first few eigenvalues. If the eigenvalues from the real data measurement were far outside the simulated distribution, I could conclude that they were not due to statistical fluctuation. Figure 5.4 shows the distributions for the first two eigenvalues. Significant correlated $1/f$ modes in the real data had much larger eigenvalues than in this simulation; therefore, the measurement was statistically significant.

To determine the importance of correlated $1/f$ modes, we computed linear combinations of the diode TODs that enhanced the $1/f$ noise. We selected the linear combinations by

diagonalizing the $1/f$ noise covariance matrix for a single CES. This covariance matrix was formed from the full correlation matrix (not a model); we weighted the diodes to have the same white noise level. Since the $1/f$ noise was frequency-dependent, we evaluated the covariance matrix at the scan frequency. The first eigenvector of this matrix corresponded to a linear combination of diode TODs with enhanced $1/f$ power at the scan frequency. We calculated 10 TODs that were linear combinations of diode TODs with coefficients taken from the first 10 eigenvectors of this covariance matrix. Since our filtering (§5.3.1) eliminated frequencies below twice the scan frequency, we needed to check the modes with power near the cutoff to ensure that our noise model accounted for the $1/f$ noise remaining after filtering. For each linear-combination TOD we computed the noise power spectrum and compared it to a noise model. The noise model included only intra-module correlations. Even so, the agreement was good above the filter cutoff frequency. Since only frequencies above the cutoff survived through the pipeline, agreement above the cutoff was sufficient for a noise model. Since the correlated component had slope 2, it was reasonable for it to be insignificant compared to the slope-1 component at high frequencies. Thus, we chose to ignore inter-module correlations in the $1/f$ noise model.

To check that inter-module correlations can be ignored for all CESes, we repeated the process in the above paragraph for 200 CESes (selected from a randomly chosen 2-week period). We computed the first 10 eigenmodes of the covariance matrix for each CES. For each eigenmode we computed the diode–linear-combination TOD. For each such TOD we computed the noise power spectrum and the χ^2 between that spectrum and the model prediction (including only intra-module correlations) between twice the scan frequency and 1 Hz. Figure 5.5 shows the distribution of these χ^2 values. Although most CESes were well-modeled with only intra-module correlations, a few CESes were not, as indicated by the large χ^2 outliers. Most of the outliers could be eliminated by including in the noise model the first inter-module mode of the correlation matrix.

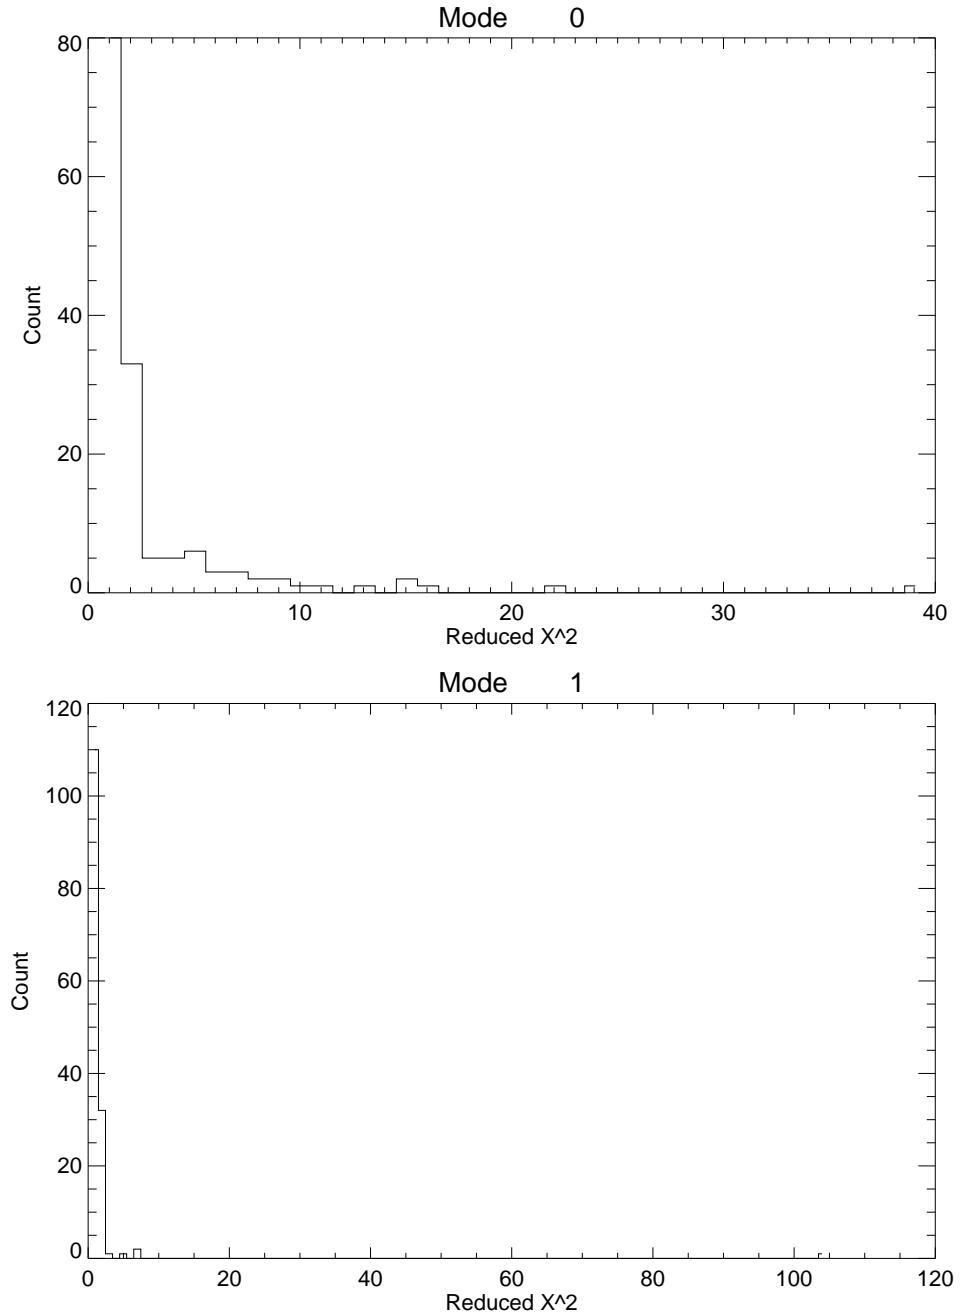


Figure 5.5: Our $1/f$ model reproduced the $1/f$ noise above twice the scan frequency. The two panels show the reduced- χ^2 histograms from the model fit to the two modes with the largest $1/f$ variance at the scan frequency. The model did not include inter-module correlations. The outliers were due to the small number of CESes where the inter-module correlations were important and scan-synchronous pickup, which cannot be modeled by a $1/f$ spectrum.

Since we chose to consider only intra-module correlations in the final noise model, we estimated the error introduced because of the few CESes that had significant inter-module correlations even above the pipeline filter cutoff. We picked seven CESes¹¹ that were outliers of this type. For these CESes, we prepared two noise models. The “array correlated” model had only the first eigenmode of the 1/f correlation matrix. In this model, we adjusted the slopes and knee frequencies for this mode in an attempt to extract the slope-2 component. From the measured slope and knee frequency of each diode, we calculated the effective knee frequencies of slope-1 and 2 components. Then we used the slope-2-component knee frequency, and the measured slope or 2, whichever was greater. We made an additional adjustment to the knee frequency

$$f' = f_0 \rho_1^{1/\alpha}, \quad (5.40)$$

where f_0 is the knee frequency before this adjustment and ρ_1 is the diagonal element of the correlation matrix constructed from the first mode. This adjustment accounted for the fact that the slope-2 component was not perfectly correlated—in the fiducial pipeline noise model we included intra-module correlations. Therefore in this validation simulation, we did not include the 1/f noise that was not correlated among modules. The other noise model, “module correlated,” used the same knee frequencies and slopes as array correlated, but we set the correlation-matrix elements between different modules to 0. The difference between “array correlated” and “module correlated” was the effect of the correlation not included in the fiducial pipeline model. At low ℓ , the difference was at most 10% of the noise bias. Since at most 10% of CESes had an important intermodule mode and the noise bias scaled linearly with the number of CESes, the effect was diluted to 1% when all CESes (including the majority that were not outliers) were included. We considered 1% to be an acceptable systematic error in the noise model¹²; therefore, we justified using only intra-

11. CESes 5410.1, 5410.2, 5420.2, 5420.3, 5420.4, 5456.2, and 5475.2

12. This noise-model error affected only the statistical error estimate, not the power-spectrum central

module correlations in the pipeline. Note that in this simulation the EE and BB noise-bias differences had opposite signs. This suggested that due to the limited number of CESes, the inter-module mode happened to couple to E and B differently. If there was no preferred orientation, we would expect this effect to average down as more CESes were included; although, we did not rely on this reduction in justifying the noise model.

5.2.2 Model of Noise Correlations

Based on the investigation above, we included only intra-module correlations in the noise model. For white noise, we simply used the noise level for each diode and correlation matrix $\hat{\rho}_{ij}^w$. For 1/f noise, we used the fitted knee frequency, slope, and correlation matrix $\hat{\rho}_{ij}^{1/f}$. To generate a noise realization from this model, we needed to diagonalize the correlation matrix. Because of the small bandwidth in the 1/f noise measurement, the matrix may be ill-conditioned. In the diagonalization, I applied the same regulation procedure described above except that diodes with knee frequencies in $(0, 10]$ mHz were not removed. Figure 5.6 shows the eigenvalues obtained for one CES. The first 2–3 modes were significant. Figure 5.7 shows the eigenvectors. The most important modes had negative correlation between the two Q diodes and between the two U diodes. We expected the – sign because the two diodes had opposite signs. I included the first three modes of the regulated 1/f correlation matrix in the noise model whenever their eigenvalues were positive. Negative eigenvalues indicated measurement failures in the correlation matrix because they correspond to unphysical negative variance modes, which cannot be modeled. The fourth eigenvalue scattered about zero, indicating that it came primarily from measurement noise. We attributed the three significant modes to MMIC-amplifier–gain fluctuations (two modes from two amplifier chains) and atmospheric fluctuations (one mode). Except for σ , the data processing was independent of the noise model. Therefore, these choices about how to model the noise correlations did not

value. Our cross-correlation method (§5.4.3) made the central value insensitive to noise bias.

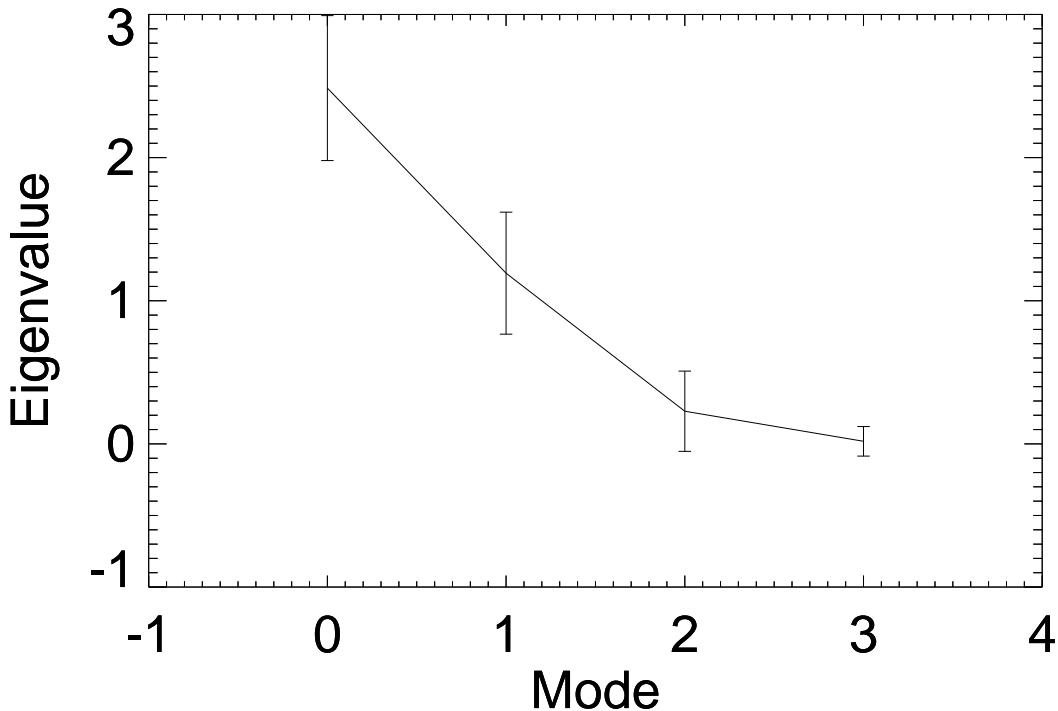


Figure 5.6: The first 2–3 1/f modes in each module were significant. The error bars show the scatter (standard deviation) among modules for one CES. Note that the last eigenvalue (0.02 ± 0.10) scatters around 0, indicating that it is dominated by measurement noise. We attributed the three significant modes to MMIC-amplifier–gain fluctuations (two modes from two amplifier chains) and atmospheric fluctuations (one mode).

bias the result. They only affected the estimate of the statistical uncertainty at the percent level.

5.3 Time-ordered–data Processing

After calibration and characterization and before map making, we filtered the TOD. First, we applied several filters to each CES-diode to remove various noise and contamination effects. Then I removed CES-diodes from further analysis if contamination remained after filtering. These steps ensured that the data used in the power-spectrum calculation had well

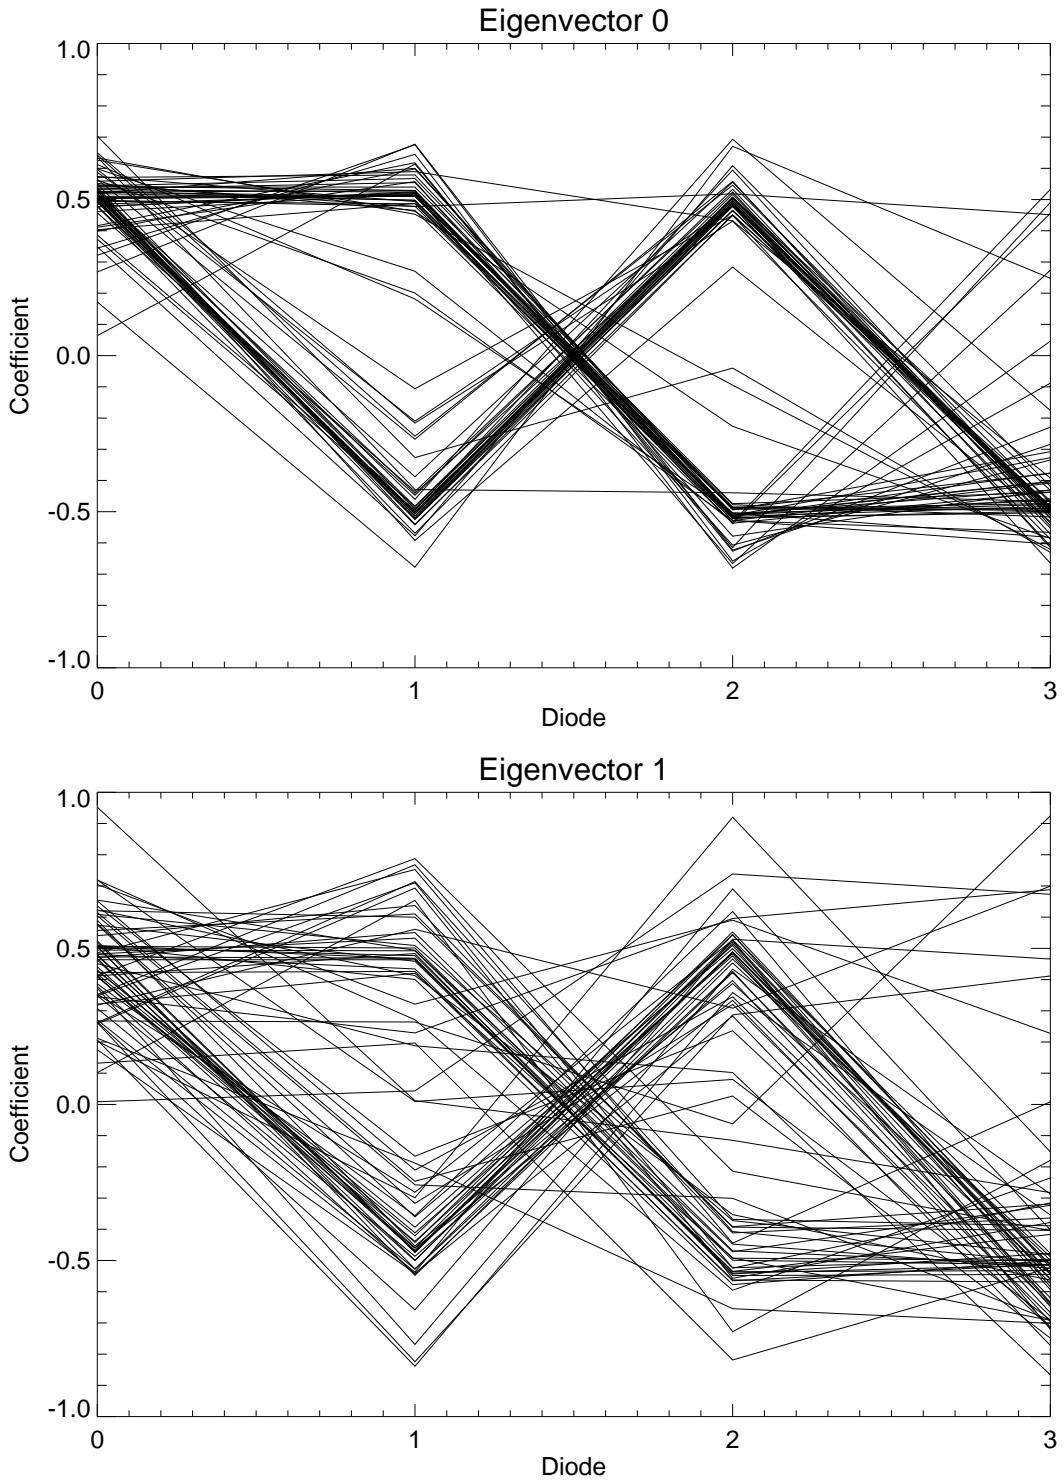


Figure 5.7: We expected opposite signs for Q1 (U1) and Q2 (U2) 1/f noise within each module. Panels show eigenvectors corresponding to the first two eigenvalues from one CES. Each line corresponds to one module. Diodes 0, 1, 2, and 3 are the same as Q1, U1, U2, and Q2, respectively.

understood noise properties.

5.3.1 Filtering

We filtered each CES-diode's TOD to remove residual 1/f noise and ground pickup. The filters were

1. We subtracted the slope (linear function of time) from each CES-diode. This prepared the data for Fourier transformation.
2. We removed high-frequency components by multiplying by a filter function in the Fourier domain. This function was unity below 9.6 Hz and zero above 9.7 Hz. Between those frequencies we smoothed with a cosine-squared function shape. We chose the filter frequency to remove a variety of high-frequency spikes (including strong contamination at 10 Hz due to aliasing of the 60-Hz power line). The signal from the maximum multipole $\ell = 1000$ entered the data at ≈ 5 Hz so the filter did not remove significant signal (see also Figure 5.13).
3. We subtracted the slope (linear function of azimuth) from each scan half-period. This filter introduced a high-pass cutoff at twice the scan frequency, thereby removing remaining 1/f noise.
4. We subtracted the average (constant during the CES) value from each datum. For this filter, we binned the data into 40 bins in azimuth. The subtracted mode was an arbitrary function of azimuth that was constant during each CES. We used this filter to remove ground pickup caused by sidelobes hitting the ground.

5.3.2 Data Selection

The goal of data selection was to remove data that, even with filtering, did not fit our noise model. I based most of the selection on the statistics calculated for each CES-diode. For

each such statistic, I determined a cut threshold based on its distribution, examination of outliers, and the validation tests (Buder, 2011d,e,f). The validation tests served as the final check that data-selection thresholds were adequate. I removed data in the following steps:

1. I removed six non-functional modules. Modules RQ7, 8, 28, and 42 had electrical connection problems, and we could not bias them. Modules RQ3 and 70 had frequent level shifts: (almost) every CES contained a large (compared to the noise level) jump in the DD level.
2. I removed four non-functional diodes. RQ4U2, RQ38Q1, and RQ40Q1 had connection problems. RQ35Q1 had a noise spike in the Fourier domain whose frequency was unstable. This problem was due to a bad ADC channel (the same problem happened in Q band; the ADC channel was the only common element.).
3. I removed RQ57Q1 when the TP (§1.3.1) was above 0.5 V. This diode had two states (Figure 5.8). The higher-voltage state had systematically worse quality (Figure 5.9).
4. I removed unusable CESes. §4.5.2 describes DAQ problems that caused an entire CES to be unusable. I also removed CESes during which the mount stalled, deck moved, generator failed, temperatures were unstable, or module bias was unusual. If a significant fraction of the data could be recovered, I redefined the CES to exclude the bad data.
5. I removed CES-diodes for which the Preamp power supply failed (Hasegawa, 2011a). The symptom of this failure was that the noise levels for one MAB increased greatly. There were six such failures during the season. We replaced the failed power supplies within a few days.
6. I removed CES-diodes when one of the sidelobes passed through the Sun (Chinone, 2011b). This created strong features in the map, and we chose selection criteria based

on pointing information to remove CES-diodes that had evidence of such a feature.

7. I removed CESes with poor cryostat temperature regulation. I excluded all CESes for which the mean temperature (at RQ9) was $> 30\text{ K}$ or the RMS fluctuation during the CES was $> 0.19\text{ K}$.
8. I removed CES-diodes whose knee frequencies were greater than twice the scan frequency. Such data had significant $1/f$ noise above the filter cutoff, often had bad weather, and made the modeling of $1/f$ noise too important.
9. I removed very short CESes with duration $< 1000\text{ s}$.
10. I removed data with bad weather. This cut removed entire CESes based on the array-average fluctuation of TP ([Kusaka, 2011c](#)).
11. I removed CES-diodes with TOD glitches. I identified glitches by calculating the maximum excursion (compared to the noise level) on three time scales: 20 ms (1 sample), 100 ms, and 1 s. If any glitch was $> 6\sigma$, I removed the CES-diode.
12. I removed CES-diodes for which the dependence on azimuth (slopes removed in filtering) had a large change within the CES.
13. I removed CES-diodes with poor fitting of the noise model. If the reduced χ^2 of the fit was > 3 below 1 Hz or > 2 above 1 Hz, I removed the CES-diode. (Because we fit the noise model before filtering but computed this statistic after filtering, we had to compute the effect of the filter on the expected noise ([Kusaka, 2010a](#))).
14. I removed CES-diodes whose noise-model fit χ^2 was > 5 in the 40 modes closest to the scan frequency. For this statistic, I applied only the ground-subtraction filter.
15. I removed CES-diodes whose Type-B χ^2 was > 10 . I calculated this statistic as the reduced χ^2 from a linear fit to the unfiltered DD vs. TP relation (Figure [5.10](#)).

16. I removed CES-diodes with level shifts during the CES. We calculated the maximum level shift on three time scales. If the shift was $> 6.5\sigma$ in 1s, $> 9\sigma$ in 10s, or $> 25\sigma$ in 100s, I discarded the CES-diode.
17. I removed CES-diodes with narrow spikes in the noise power spectrum. I averaged the noise spectrum in logarithmically-spaced frequency bins. If any bin was $> 5\sigma$ above the noise model, I discarded the CES-diode. Because I found strong contamination near 1.2 Hz from the cryogenic refrigerator cycle, I removed CES-diodes with any bin between 1.15 and 1.3 Hz $> 4\sigma$. This cut also removed spikes due to scan frequency harmonics (Figure 5.11).
18. I found CES-diodes whose noise level changed during the CES (Figure 5.12). I calculated a noise non-stationarity statistic as the standard deviation of the standard deviations of 10-s periods. If this statistic was > 2.5 times the expected variation, I removed the CES-diode.
19. I removed CES-diodes with very large scan-synchronous signal. Although we filtered the scan-synchronous signal with the ground subtraction filter, we found evidence of residual contamination (Nagai, 2012a). Therefore I removed any CES-diode whose noise power at the scan frequency was > 100 times the expected level.

Table 5.3 lists the selection efficiency for each patch.

5.4 Power-spectrum Calculation

The power-spectrum calculation was a multi-step process. First, we made sky maps from the filtered TOD. Then we computed spectra with a pseudo- C_ℓ estimator. In this process, we divided the maps into cross-correlation units and subtracted the spectra from each unit so the result depended only on cross-correlations. Then we divided by transfer and window functions to account for the reduction of power from filtering and smoothing by the beam.

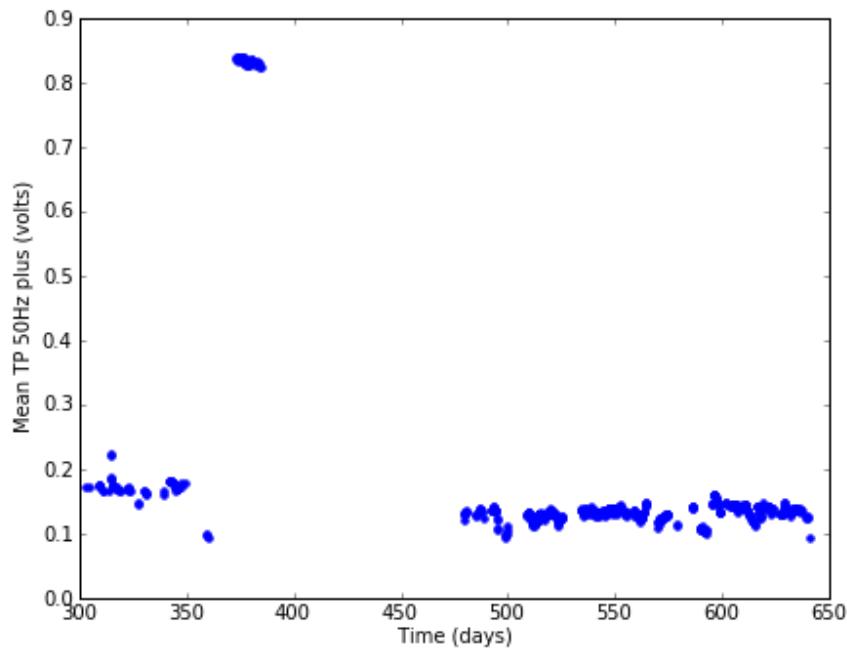


Figure 5.8: RQ57Q1 had two TP states. I rejected data from it in the high state because the quality was bad.

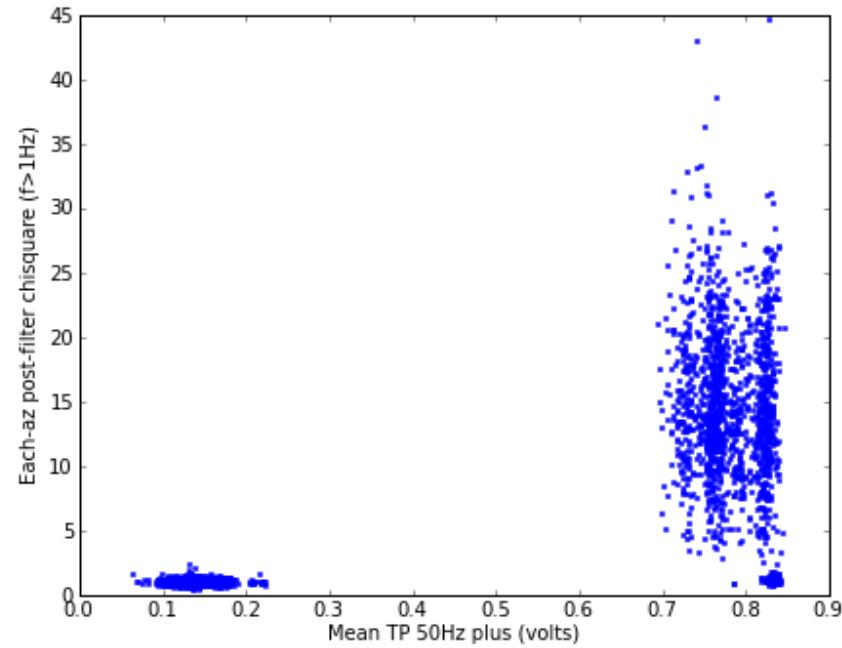


Figure 5.9: In the high-TP state, RQ57Q1 had very bad noise-model fit χ^2 .

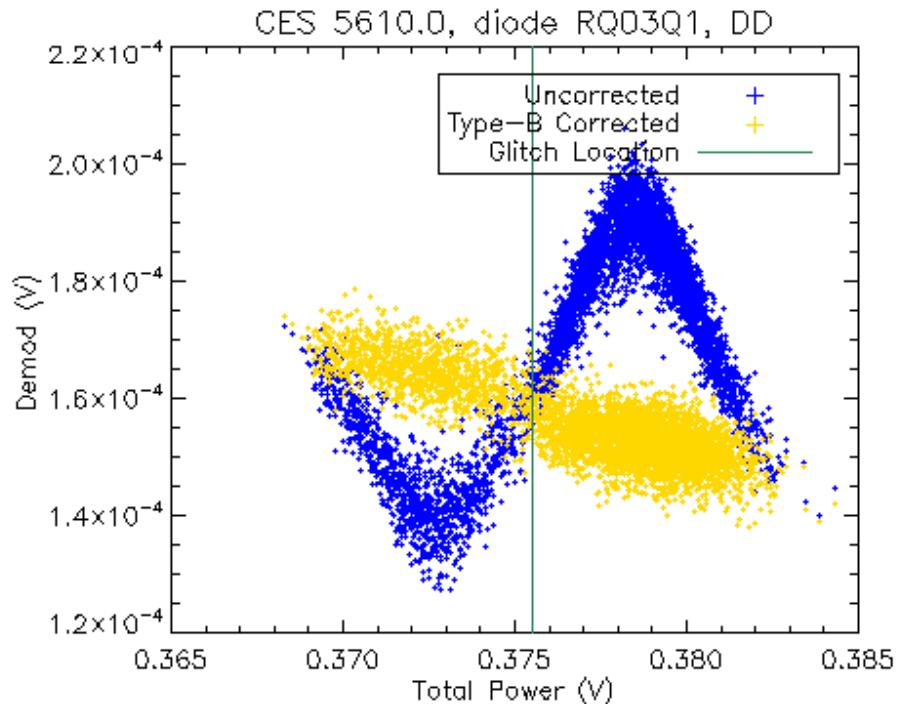


Figure 5.10: I rejected CES-diodes with large χ^2 in a linear fit to the DD vs. TP relation. Type-B glitching created non-linearity in this relation.

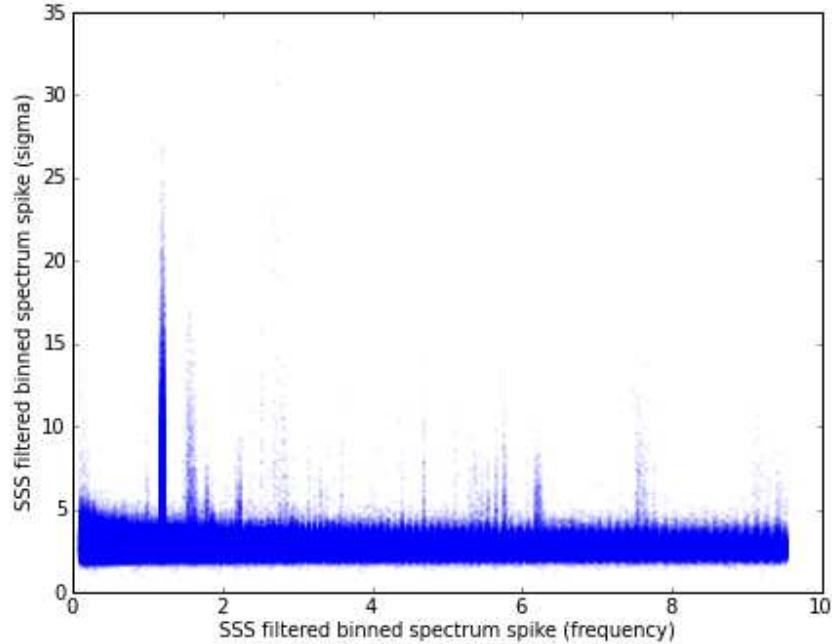


Figure 5.11: I rejected CES-diodes with a noise-spectrum–spike statistic $> 5\sigma$. The strongest spikes at 1.2 Hz were due to the refrigerator cycle.

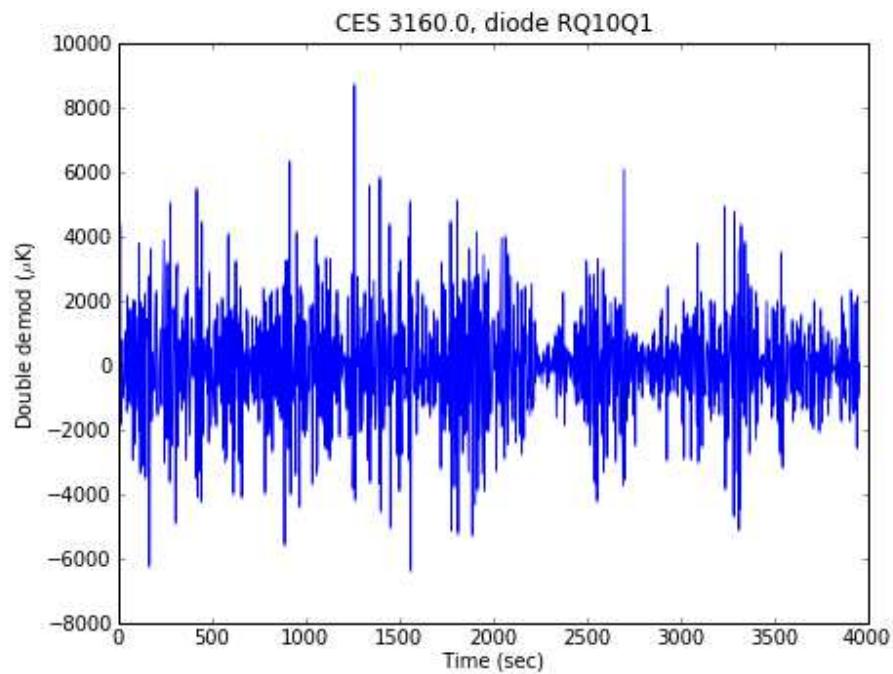


Figure 5.12: I rejected CES-diodes whose noise level changed within the CES.

Table 5.3. Data-selection Efficiency

Patch	Baseline CES-diodes ^a	Selected CES-diodes	Efficiency
2a	448607	286643	64%
4a	364443	247448	68%
6a	298943	176107	59%
7b	173843	113038	65%
Total	1285836	823236	64%

^aTo compute the baseline I applied only cuts 1, 2, 3, 4, 5, 7, and 9. These corresponded to instrument or observing failures rather than analysis limitations that required a certain data quality.

5.4.1 Map Making

We made a naive map from each CES-diode. For each TOD sample, we estimated the polarization

$$\begin{pmatrix} \hat{Q} \\ \hat{U} \end{pmatrix} = \left(\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P} \right)^{-1} \mathbf{P}^T \mathbf{N}^{-1} \vec{d}, \quad (5.41)$$

where

$$\mathbf{P} \equiv \begin{pmatrix} \cos 2\psi & \sin 2\psi \end{pmatrix} \quad (5.42)$$

specifies the polarization angle (ψ) of the diode,

$$\mathbf{N} \equiv \begin{pmatrix} \sigma^2 \end{pmatrix} \quad (5.43)$$

is the variance (including only the white noise and ignoring correlations), and \vec{d} is the sample. We assigned each sample to a pixel using the pointing model. The maps used the HEALPix¹³ (Górski et al., 2005) pixelization scheme with the parameter $N_{\text{side}} = 512$ corresponding to $0^{\circ}1$ pixel resolution. We combined multiple observations of the same pixel with inverse-variance weighting. This naive combination was not optimal; however, the statistical uncertainty increase was percent-level (Kusaka, 2009). We used the same weighting to combine different diodes and CESes.

5.4.2 Pseudo- C_ℓ Calculation

From combined maps of many CES-diodes we estimated the power spectra. We used the MASTER (Monte Carlo Apodized Spherical Transform Estimator) technique to calculate an unbiased estimator from partial-sky data with filtering (Hivon et al., 2002; Hansen & Górska,

13. <http://healpix.jpl.nasa.gov>

2003). In this technique, we first calculated “pseudo power spectra”

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle, \quad (5.44)$$

where $M_{\ell\ell'}$ is a mode-coupling matrix¹⁴ that accounts for the limited sky fraction, $F_{\ell'}$ is the transfer function that accounts for filtering, and $B_{\ell'}$ is the beam window function that accounts for the instrument’s resolution. The brackets indicate the relation between the true and pseudo spectra holds on average for an ensemble of sky realizations. We converted the pseudo spectra back into true spectra estimates after binning in ℓ to create 19 bandpowers

$$\hat{C}_b = \sum_{b',\ell} F_b^{-1} K_{bb'}^{-1} P_{b'\ell} \tilde{C}_\ell, \quad (5.45)$$

where F_b^{-1} is the (binned, see §5.5) inverse transfer function,

$$K_{bb'} = \sum_{\ell,\ell'} P_{b\ell} M_{\ell\ell'} B_\ell^2 Q_{\ell'b'} \quad (5.46)$$

accounts for the mode-mode coupling and beam, and $P_{b\ell}$ and $Q_{\ell'b'}$ are the binning projection and interpolation functions. We verified that \hat{C}_b was unbiased with simulations (§5.5).

We calculated B_ℓ from the calibrated beam shape (§5.1.1 and Monsalve (2012a,b)).

$$B_\ell = \frac{2\pi}{\Omega} \int b(\theta) P_\ell(\cos \theta) d(\cos \theta), \quad (5.47)$$

where Ω is the solid angle¹⁵, $b(\theta)$ is the beam shape, and P_ℓ are Legendre polynomials. We multiplied B_ℓ calculated from the beam shape by an additional factor corresponding to a Gaussian beam with FWHM = 5'1 to account for the pointing-model scatter (§5.1.4).

14. This estimator, combined with the limited sky fraction, mixes E and B modes. The coupling matrix includes this mixing as well; although I suppressed the indices for different polarization modes.

15. Some conventions differ by not including Ω .

When we calculated power spectra, we masked known point sources. In the MASTER technique we weighted pixels based on how often we observed them (Feldman et al., 1994). To avoid contamination from point sources, we reduced the weight to zero for nearby pixels. We masked pixels within 2° of Centaurus A (in patch 2a) and within 1° of Pictor A (in patch 4a). The mask was not apodized i.e. outside those radii the pixel weights were unaffected.

5.4.3 Cross-correlation

We used a cross-correlation technique to eliminate noise bias and reduce systematic errors. We made separate maps for each pointing division. Each division included a range of azimuth and deck angles chosen to evenly divide the data. There were 40 such divisions: 5 in azimuth¹⁶ and 8 in deck; although, patches 6a and 7b did not have data in all divisions¹⁷. We performed the pseudo- C_ℓ calculation for each pointing division as well as the cumulative map (sum of all divisions). Then we subtracted the power-spectra estimates of each division from the cumulative spectra. Thus, only cross-correlations of different pointing maps contributed to the result. Because the data in different divisions were taken at different times, the noise between divisions was uncorrelated. Therefore, there was no noise bias in our estimator¹⁸. Moreover, systematic contamination that did not correlate between divisions was suppressed.

16. I made the azimuth divisions at $130^\circ 157$, $157^\circ 52$, $209^\circ 935$, and $232^\circ 254$ East of North. I chose these points to equalize the number of CES-diodes in each division.

17. Patch 6a had a lower Declination and, therefore, a smaller azimuth range than the other patches (24 total divisions). We only observed patch 7b while it was rising, limiting the azimuth range (30 total divisions).

18. Power-spectrum estimators that include auto-correlations are subject to “noise bias”—the power of the noise correlating with itself. By eliminating noise bias, we removed the possible systematic error from misestimating it.

5.5 Simulations

A simulation suite is a necessary component of the MASTER technique. Our suite contained three types of simulations: calibrating, fiducial, and noise. I based the simulations on the instrument noise model developed above. I used the simulations to calculate F_b and estimate the statistical uncertainty. The simulations also verified that the analysis method was unbiased.

The three types of simulations differed in whether I input CMB power and instrumental noise. In calibrating simulations I input CMB power but no instrumental noise. For these simulations, I assumed a flat (i.e. ℓ -independent) input power for both E and B modes. In fiducial simulations I input both CMB power and instrumental noise. Usually fiducial simulations had CMB power input drawn from realizations of a Λ CDM model; however, some simulations had non-zero B-mode power to estimate the sample variance for B modes. Noise simulations had instrumental noise only. For the final configuration I performed 200 calibrating, 200 fiducial, and 100 noise simulations. Each simulation was a complete realization of the experiment observation and TOD.

The simulations included white noise and 1/f noise. Both noise sources were correlated within each module but not between modules. For white noise, I added random Gaussian values with correlation matrix $\hat{\rho}^w$. For 1/f noise, I created random Gaussian values in the Fourier domain with correlation matrix $\hat{\rho}^{1/f}$ after regulation described in §5.2.2. Then I scaled the noise to have the measured 1/f slope and knee frequency. I generated more samples than the data length of the CES so the 1/f noise at the beginning and end were uncorrelated.

I computed F_b from the calibrating simulations. The ratio between the simulation input and output spectra is F_b (Figure 5.13). Since there was no EB power in the simulations, I computed the EB transfer function as the geometric mean of the E and B-mode functions.

The fiducial simulations confirmed that computing the binned transfer function was sufficient; otherwise I would need the transfer function at every ℓ . This sufficiency depends on the input spectra. If I had found a result radically different from Λ CDM, I would have had to re-evaluate whether the binned transfer function would still provide an unbiased estimator. We also confirmed that the filtering did not mix different ℓ ([Kusaka, 2008](#)); otherwise the transfer function would need to be a matrix including that mixing possibility. (An additional complication in computing the transfer function is that simulated skies have finite resolution and pixel size, but the true sky has infinite resolution. Normally the simulation input was pixelized at $N_{\text{side}} = 1024$. To solve this problem I ran additional calibrating simulations with $N_{\text{side}} = 2048$. We also accounted for the finite resolution of TOD samples ([Kusaka, 2012b](#)).)

The fiducial simulations confirmed the pipeline did not introduce bias. If I input a Λ CDM spectrum, I recovered the same spectrum ([Figure 5.14](#)). The simulations also provided the primary estimate of the statistical uncertainty. Unless otherwise noted, power-spectra error bars are the standard deviations of the fiducial simulations.

5.6 Validation Tests

We performed a set of validation tests to confirm that the analysis was correct. These tests were part of a blind-analysis strategy wherein we finalized the analysis method without knowing the results. The most important validation test was the “null suite,” a set of tests for consistency between two halves of the data. As a second validation test, I checked that the spectra computed from different patches were consistent. Finally, I compared the results of different analysis configurations.

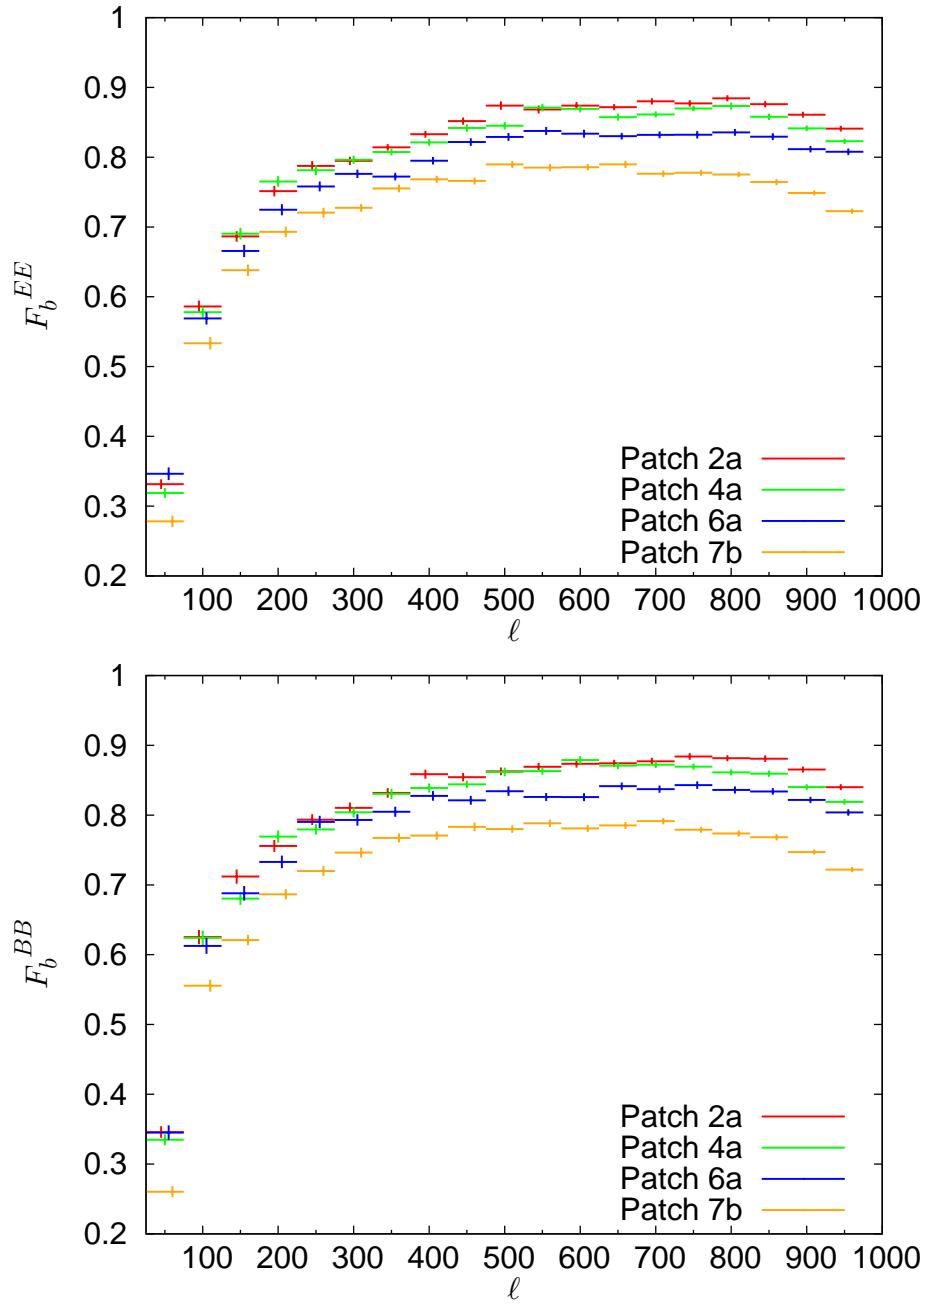


Figure 5.13: The transfer function accounted for filtering in the pipeline. I computed the transfer function from noiseless simulations; the error bars indicate statistical uncertainty from the 200 simulations of each patch. Except at low ℓ , filtering did not affect the CMB power by more than 20%.

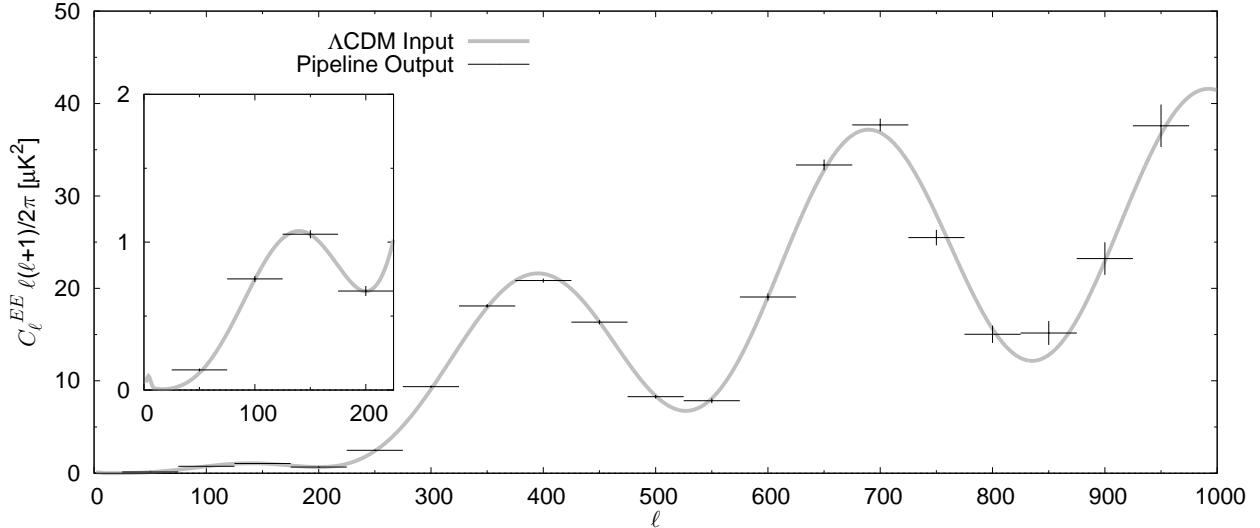


Figure 5.14: Simulations confirmed the analysis pipeline was unbiased. Error bars indicate statistical uncertainty from the 197 full-season simulations of patch 2a.

5.6.1 Blind-analysis Strategy

Blind analysis removes opportunities for observer bias by making the analysis choices independent of our knowledge of the result (Klein & Roodman, 2005). There were many choices to make in the analysis method. If we calculated the results from all possible choices, we could choose the result closest to our a priori beliefs. These beliefs might include a preference for results close to Λ CDM or for the lowest B-mode power. Even if we did not intentionally alter the analysis method to prefer such results, there are more subtle effects such as “stopping bias” that can cause results to conform to previous experiments and predictions¹⁹. Collectively such effects constitute observer bias, a systematic error that is difficult to quantify because it depends on the psychology of physicists. By performing a blind analysis—one in which we did not see the result until after making all analysis choices—we removed this systematic error.

The method of blind analysis we chose for QUIET was not to know the result (i.e. the

19. Croft & Dailey (2011) provide some evidence that cosmological-parameter measurements have observer bias.

Table 5.4. Null-suite Probabilities to Exceed

Patch	χ bias	Largest $(\chi_b^{\text{null}})^2$	Total $(\chi_b^{\text{null}})^2$
2a	26	14	78
4a	46	40	6
6a	74	45	50
7b	8	18	76

CMB polarization power spectra) until we chose the calibration and analysis methods. We did look at TOD and maps including the CMB signal because the relationship between them and power spectra was not straightforward. Since we did not let the results inform the analysis method, we used the validation tests to determine whether analysis choices were correct. For the most important tests (the null suite and patch comparison) we set the criteria to pass the tests in advance (Buder et al., 2011). We also evaluated the systematic errors (§5.7) before seeing the result.

5.6.2 Null Suite

The null suite was the most important validation test. The suite had 32²⁰ null tests, each of which divided the data in half to target a known problem or potential systematic error. We decided in advance to determine the success of the null suite from three criteria for each patch (Table 5.4). In addition to these main criteria, we examined many other statistics and distributions to look for surprises. We found two such surprises: a bias in the results including auto-correlations and excess χ^2 at low ℓ . To confirm these effects did not indicate a failure in our noise model, we performed additional, random null tests that validated the noise model with high precision.

20. One of the null tests, based on whether the patch was rising or setting during observations, did not work for patch 7b because we always observed it while rising. There were only 31 null tests for patch 7b.

The null-test calculation followed the normal analysis method including pseudo- C_ℓ and cross-correlation. I divided the CES-diodes²¹ into two halves and computed a separate map from each half. I chose the pixel weighting

$$W(x) = \frac{1}{\sigma_1(x)^2 + \sigma_2(x)^2}, \quad (5.48)$$

where $\sigma_1(x)^2$ is the variance for pixel x after accumulating all CES-diodes in half 1. The choice of W fixed the mode-coupling matrix ($M_{\ell\ell'}$ and $K_{bb'}$, §5.4.2). Then I defined the null-power estimator

$$\hat{C}_b^{\text{null}} \equiv \frac{K_{bb'}^{-1} P_{b'\ell} \tilde{C}_\ell^{11}}{F_b^{11}} - 2 \frac{K_{bb'}^{-1} P_{b'\ell} \tilde{C}_\ell^{12}}{F_b^{12}} + \frac{K_{bb'}^{-1} P_{b'\ell} \tilde{C}_\ell^{22}}{F_b^{22}}, \quad (5.49)$$

where \tilde{C}_ℓ^{11} is the pseudo-power of map 1, \tilde{C}_ℓ^{22} is the pseudo-power of map 2, \tilde{C}_ℓ^{12} is the pseudo-power from the correlation between maps 1 and 2, and F_b^{ij} are the corresponding transfer functions²². The expected value of \hat{C}_b^{null} was zero. Because the two null-test halves had different filtering and scanning, the difference map is not necessarily null²³. \hat{C}_b^{null} is the unique null-power estimator (Smith, 2009) that

1. is invariant if the maps are rescaled by constants (during filtering)
2. is invariant if the two halves are interchanged
3. is independent of the (CMB) power spectra

I chose the null tests to be sensitive to known instrumental problems, systematic-error sources, and noise-model imperfections. I required that no null test be $> 40\%$ correlated with

21. The two scan null tests divided each CES-diode into left-going/right-going and accelerating/decelerating halves.

22. I simulated all the null tests and computed the transfer functions from the null-test calibrating simulations.

23. If and only if all three transfer functions are identical is the difference map a null map.

any other (Figure 5.15). A large number of tests can detect contamination at a level well below the statistical threshold of a single null test. However, performing a very large number of null tests dilutes the well-motivated null tests. As the number of null tests increases, so does the largest expected statistical fluctuation. Since any outliers must be compared with the expected fluctuation, a very large number of null tests implies that any single null test has reduced discriminating ability: the level of contamination required to cause a statistically significant outlier increases with the number of null tests. Therefore the choice of the number of null tests represents a compromise between a lower detection threshold for the most well-motivated null tests and a lower detection threshold for contamination that appears in all null tests²⁴. Since in practice the number of null tests was limited by computational resources, I compromised by choosing the ≈ 30 best motivated null tests. Appendix G lists the null tests.

I constructed the null-suite passing criteria from

$$\chi_b^{\text{null}} \equiv \hat{C}_b^{\text{null}} / \sigma_b^{\text{null}}, \quad (5.50)$$

where σ_b^{null} is the RMS of \hat{C}_b^{null} from fiducial simulations. Based on the experience in Q band, I computed three summary statistics that revealed problems most often:

1. The average of χ_b^{null} for all null tests and ℓ bins. This statistic is called “ χ bias.”
2. The largest $(\chi_b^{\text{null}})^2$ for all null tests and ℓ bins
3. The total $(\chi_b^{\text{null}})^2$ for all null tests and ℓ bins

Each statistic included only EE and BB null spectra. Including EB was not an independent test, and EB contamination required correlation contamination in both EE and BB. (Although I did not include EB spectra in these summary statistics, I still checked that the EB

24. The random null tests provided a strong, independent test for such contamination.

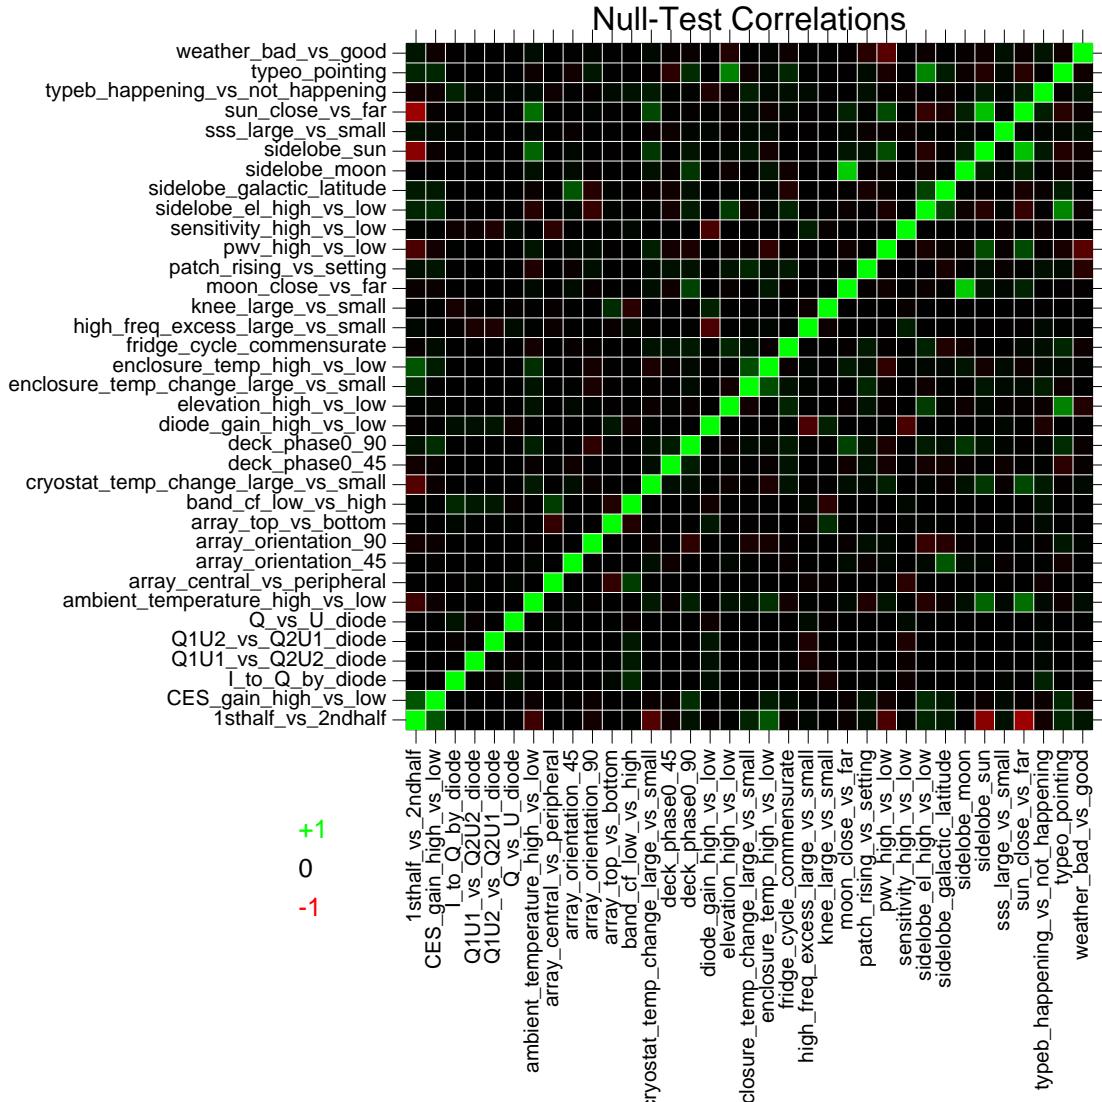


Figure 5.15: I chose null tests that were uncorrelated. Green and red indicate correlated null tests; black indicates no correlation. Appendix G defines each null test and the correlation calculation. The correlations shown are for patch 2a.

null spectra were zero-consistent.) To pass, I required each statistic have a probability to exceed (PTE) $> 4\%$. I calculated each PTE by comparing the data statistic to the distribution of the same statistics in fiducial simulations²⁵. Each patch had an independent null suite (because the patches did not overlap). For three independent tests with 4% probability to fail by statistical fluctuation, the combined failure probability generated by statistical fluctuations was 12%. However, the three tests for each patch were not independent. Therefore, the probability for a patch to fail the null suite by statistical fluctuation was $\lesssim 10\%$, which I considered reasonable. Even if the null suite passes I am not required to stop the analysis. I can examine other statistics (e.g. the worst single null test or EE alone), and we should not finish the analysis if we believe the result will be incorrect. Predefining the null-suite evaluation criteria does not force us to report an incorrect result if we find a problem outside those criteria. In the final configuration, all patches passed their null suites (Table 5.4). Moreover, the full χ_b^{null} distribution showed how well the simulations modeled the data (Figure 5.16).

In addition to the main statistics above, I examined many statistics and distributions from the null suite to check for unexpected errors. These statistics included χ^{null} from EE or BB only, the largest total $(\chi^{\text{null}})^2$ for any single null test, the RMS of the χ^{null} distribution, Kolmogorov-Smirnov tests, and χ_b^{null} as a function of b . Among all of them I found only two problematic ones: the auto-correlation bias and the low ℓ total χ^2 .

In addition to the normal analysis using cross-correlation (§5.4.3), I also computed the power spectra including the auto-correlations. Including the auto-correlations makes the result sensitive to noise bias. I subtracted the noise bias estimated from noise simulations; however, the result remained sensitive to errors in the noise model. Therefore, the auto-correlation results were an important test of the noise model. During null-suite evaluation, I computed the difference between the auto-correlation and cross-correlation null tests. I found that, on average, auto-correlation null tests had higher power, especially at low ℓ

25. I defined the χ bias PTE to be the fraction of fiducial simulations with difference from zero larger than the data.

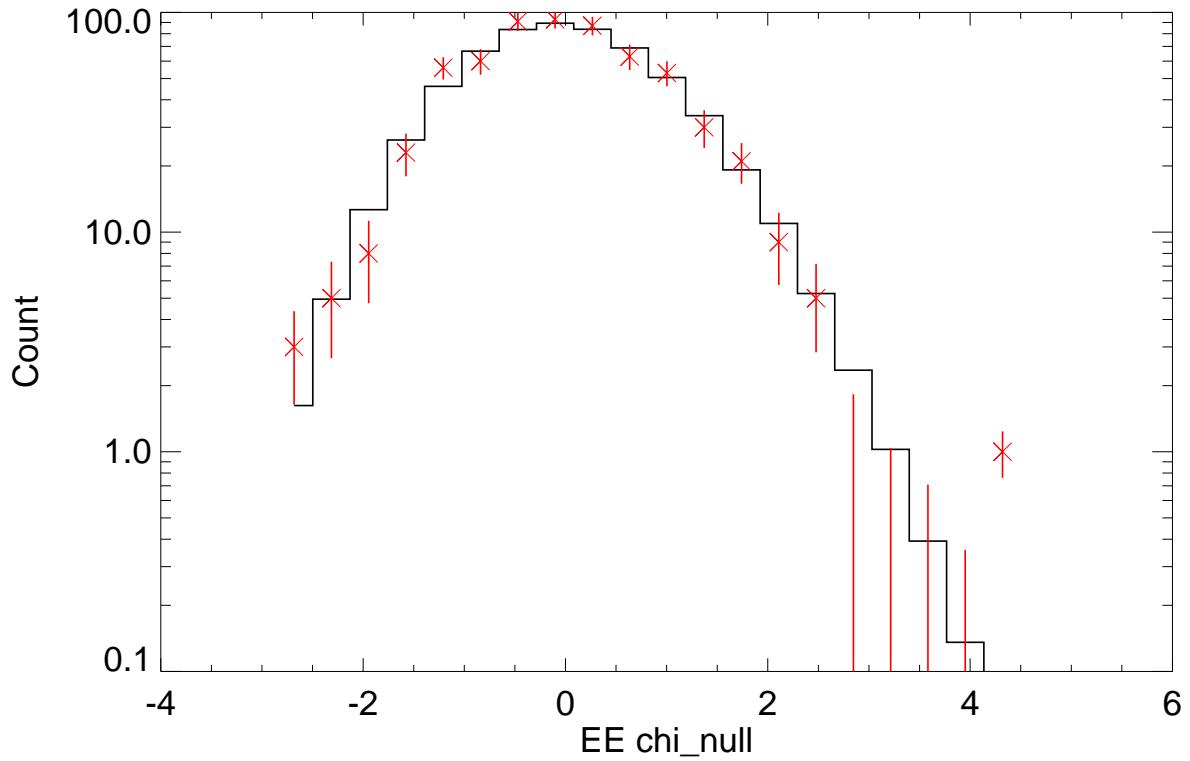


Figure 5.16: I compared the distribution of χ_b^{null} (red) to simulations (black) to confirm the noise model and limit systematic errors. The histogram shows all χ_b^{null} values for all EE bins b for all null tests of patch 2a. The error bars are the RMS fluctuations of the histogram from simulations.

(Figure 5.17). This excess indicated systematic error in the auto-correlation results or an error in the noise model. Because the final results used cross-correlation, systematic errors in the auto-correlation results were not concerning. However, an error in the noise model would cause the statistical error to be misestimated. Therefore I investigated this excess power in auto-correlation null tests (Buder, 2011b). I found that the excess was present in all null tests. I checked that the excess was not correlated with scan-synchronous signal. I found that 1/f-noise misestimation was unlikely to be the cause. Patch 2a (§4.2) had the largest excess; if the excess was due to noise misestimation it should be similar for all patches. I was unable to identify the cause of the excess.

When I totaled $(\chi_b^{\text{null}})^2$ for each bin b , I found an excess in the lowest bin ($\ell = 25$ –75) for EE. The excess was strongest in patch 4a with total $\chi^2 = 71$ for 32 degrees of freedom. Because the distribution of \hat{C}^{null} was not Gaussian, I calculated the probability of such an excess from the fiducial simulations, finding $< 0.5\%$. However, that was not the proper a priori probability. We did not decide in advance to focus on the lowest bin, and the probability of finding one outlier among many bins is much higher²⁶. The excess was smaller in patches 6a (22%) and 7b (14%), and not present in 2a (53%).

Because of the surprising statistics above, we used random null tests to confirm the noise model (Ishidoshiro, 2012). We created 1000 null tests, each of which divided the CESes randomly into two equal groups. We compared the resulting null spectra to simulations. The results confirmed our noise model was accurate to 3.1% (in C_ℓ).

5.6.3 Patch Comparison

The second most important validation test was the patch comparison. In this test, I confirmed that the four patches produced consistent power spectra. I computed the consistency

26. There is no exact probability without deciding which bins to examine, which we did not do.

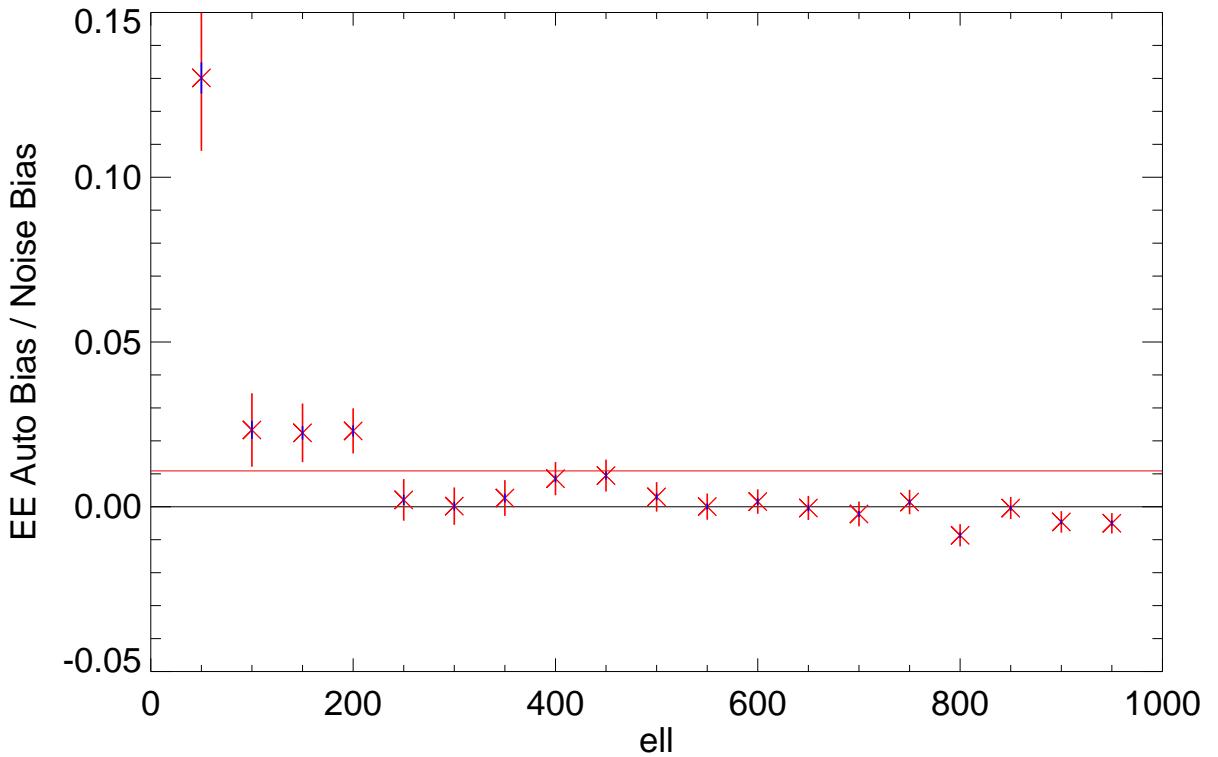


Figure 5.17: The null suite indicated excess noise in power spectra including auto-correlations. The “auto bias” is the difference between null spectra with and without auto-correlations; it is independent of the cross-correlation contribution. Dividing each bin by the noise bias yields the fractional excess noise as a function of ℓ . Our result was free from this excess because it included only cross-correlations. The plot averages all null tests for patch 2a. The horizontal red line is the average of all ℓ . The red error bars indicate statistical uncertainty; the blue error bars indicate uncertainty from noise modeling due to the finite number of noise simulations.

statistic (Buder, 2010b)

$$\chi_p^2 \equiv \sum_{i=0}^3 \sum_b \left(\frac{\hat{C}_b^i - \mu_b}{\sigma_b^i} \right)^2, \quad (5.51)$$

where \hat{C}_b^i is the bandpower²⁷ of patch i in bin b ,

$$\mu_b \equiv \frac{\sum_i \hat{C}_b^i w_b^i}{\sum_i w_b^i} \quad (5.52)$$

is the average of the four patches,

$$w_b^i \equiv 1/(\sigma_b^i)^2 \quad (5.53)$$

are patch weights, and σ_b^i are the fluctuations of each bandpower from fiducial simulations. I compared χ_p^2 to the expected distribution from fiducial simulations (Figure 5.18). Before performing the test, I set passing criteria of 4% PTE for each of EE, BB, and EB. The final PTEs were 62%, 5%, and 16% for EE, BB, and EB, respectively.

5.6.4 Configuration Comparison

I tested that the power spectra were independent of the details of the analysis choices. In this test I computed the difference of spectra from two different configurations of the analysis²⁸. Because it was computationally prohibitive to run simulations for this test, I developed a heuristic method to estimate the expected fluctuations (Appendix I). Usually I compared two data-selection choices with slightly different thresholds. In such a test the expected fluctuation was well below the statistical error on the final power spectra (Figure 5.19). Moreover, by changing the selection thresholds I directly probed the most contaminated data and confirmed that the result was insensitive to the details of data selection. I limited such

27. For computational convenience, I used the same transfer function for all patches in this test. The difference was < 10%, and I used simulations to check that there was no resulting bias.

28. To avoid possible observer bias, I randomized the sign of the difference.

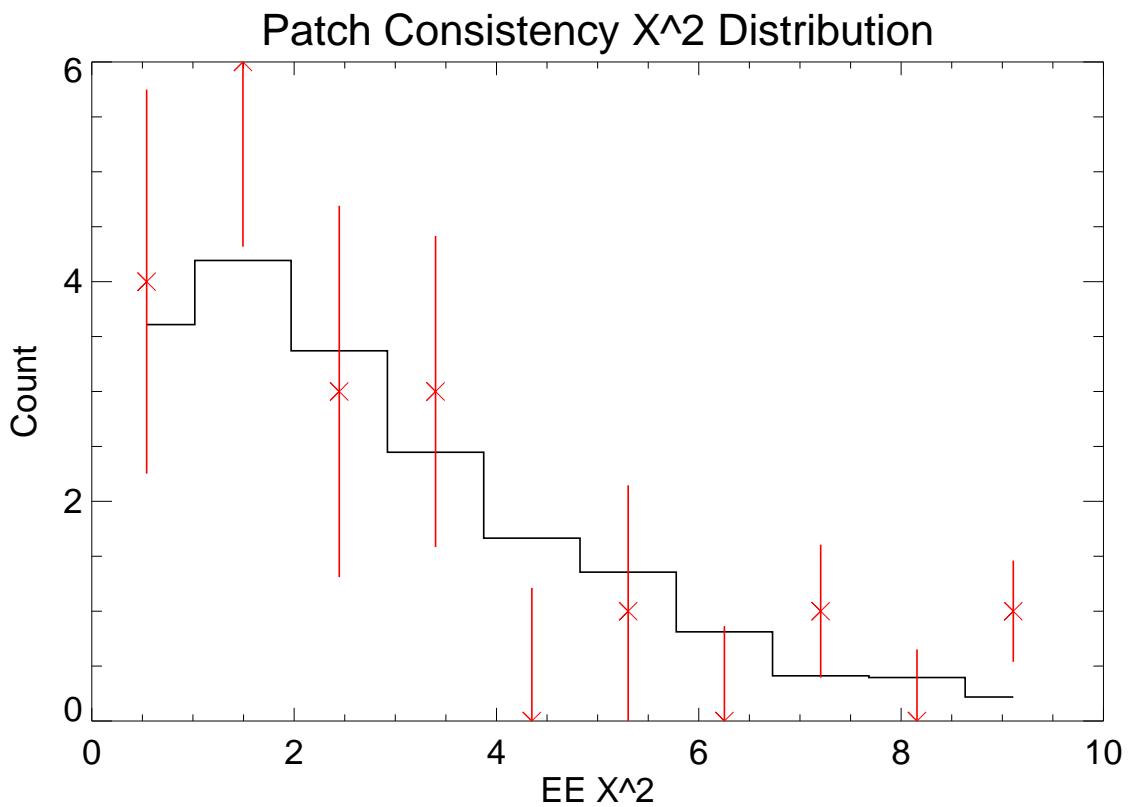


Figure 5.18: The patch comparison validation test confirmed that the four patches had consistent spectra. The χ_p^2 distribution from data (red) matched the expectation from simulations (black).

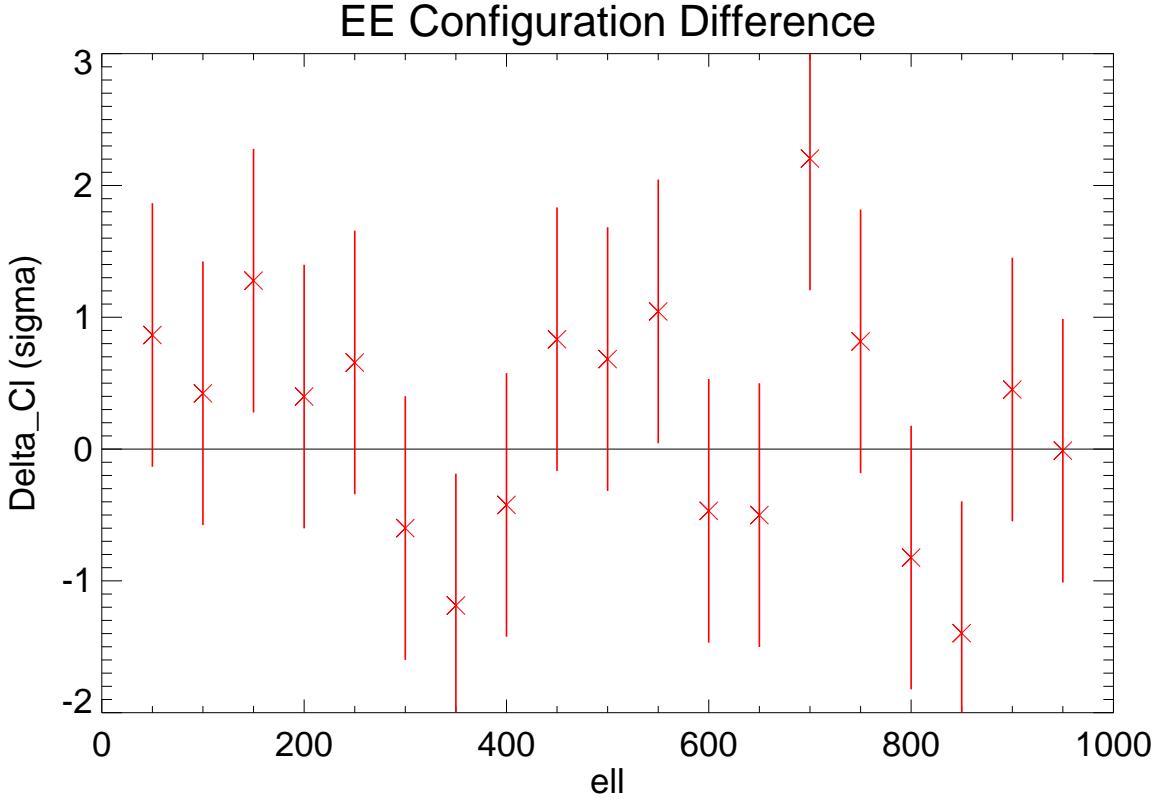


Figure 5.19: Comparison of different data-selection thresholds confirmed no contaminated data remained. The plot shows the difference between two EE spectra from patch 2a, computed with slightly different thresholds for the noise-spectrum–spike cut (§5.3.2). The uncertainty on the difference (error bars) is 17% of the total statistical uncertainty for one of the configurations.

effects to $\approx 30\%$ of the total statistical error. I used the same method to test for systematic errors by changing the configuration in ways that would be sensitive to the systematic error sources. I tested for errors in the pointing model, in ground filtering (§5.7.6), and in source masking (§5.7.6).

5.7 Systematic Errors

We supplemented the validation tests with systematic-error estimates. We used several techniques to test for systematic errors well below the statistical limits of the validation

tests. We investigated systematic errors from many possible effects:

1. Several important systematic effects were due to miscalibration—the difference between the actual and assumed calibration parameters.
2. We estimated the effect of spurious polarization created by instrumental imperfections.
3. Errors in Type-B-glitch correction created an additional contribution to this instrumental polarization.
4. We evaluated the impact of Sun contamination through the telescope sidelobes.
5. We discovered and estimated the impact of a type of polarized scan-synchronous signal.
6. Finally, we investigated many other effects and limited the corresponding errors to well below both the statistical uncertainty and the important systematic errors above.

For the effects that were significant, we assigned a corresponding systematic error to the power-spectra results. Systematic errors were the dominant source of uncertainty in the E-mode spectrum. We paid special attention to the $\ell \approx 100$ region of the B-mode spectrum because that is where we are searching for evidence of inflation. For the B-mode spectrum, we controlled systematic errors to below the level of $r = 0.01$ in the critical $\ell \approx 100$ region (Figure 5.20).

5.7.1 Miscalibration

The differences between the true calibration parameters and the ones we used in analysis (§5.1) created systematic errors. Our strategy to address this miscalibration was to compare several alternate calibration methods, and use the difference between them to estimate the possible error in calibration. The important calibration parameters were the beam shape, responsivity, detector angles, and pointing model.

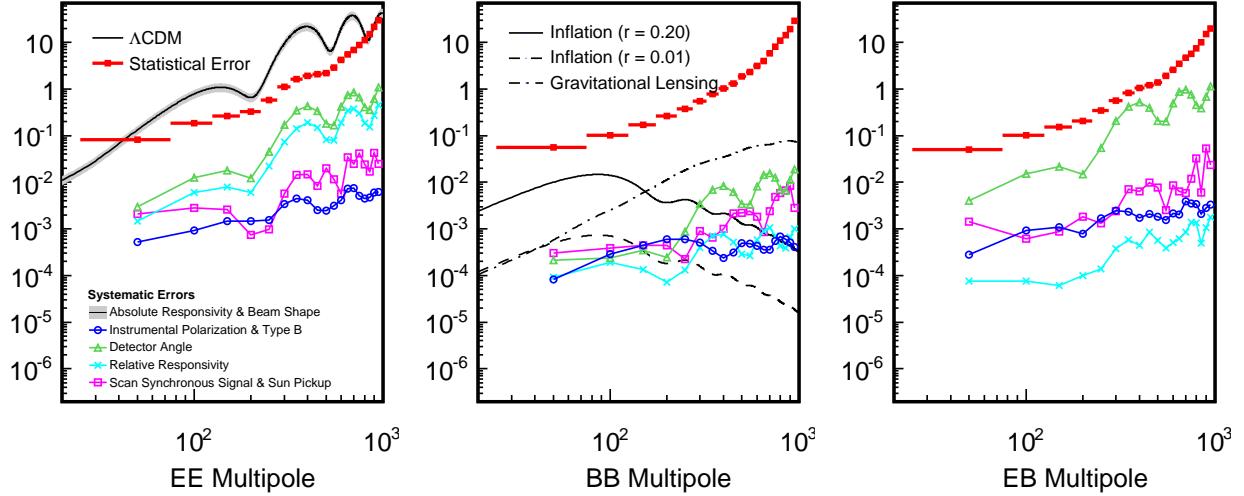


Figure 5.20: We considered many possible systematic errors. Most were negligible compared to the statistical errors; however, beam-shape and responsivity miscalibration were dominant for C_{ℓ}^{EE} . All systematic errors were below the level of $r = 0.01$ for $\ell < 125$ B modes. Units are $C_{\ell\ell}(\ell + 1)/2\pi$ in μK^2 . SOURCE: Osamu Tajima

To estimate the systematic errors from beam-shape miscalibration, we compared the fiducial calibration to optics simulations and pre-deployment measurements. We calculated the difference in the beam size (FWHM) when we used optics simulations instead of the Tau A measurements. We calculated the difference in the beam shape when we used differential-temperature measurements of Jupiter instead of the antenna-range measurements. We propagated these differences to the window function and assigned them as systematic errors. Because the window function is a multiplicative factor, these errors are proportional to the power spectra. We used the same method to assign the error (4.8%) for the beam solid angle ($13.63 \pm 0.65 \mu\text{sr}$, §5.1.1).

To estimate the systematic errors for responsivity miscalibration, we propagated the responsivity model uncertainties (§5.1.2) to the systematic errors in the power spectra. The absolute responsivity uncertainty (8.3%, 17% in C_{ℓ}) created a multiplicative error independent of ℓ . The power spectra are proportional to the responsivity squared so the responsivity uncertainty is double in the spectra. For the uncertainty on relative responsivity among

diodes, we used simulations to compare the effects of different models on the spectra. The largest effects came from changing the model to use only Tau A measurements. We also calculated the effects of fluctuations (which we modeled as random) in the relative responsivity. Relative responsivity errors are not purely multiplicative and created small, but not negligible, B-mode power (Nagai, 2012e).

To estimate the systematic errors from detector-angle miscalibration, we ran simulations with alternate detector-angle-calibration methods (§5.1.3 and Nagai (2012b)). For the absolute calibration error (shift in the average angle), we ran a simulation with all the detector angles shifted by the calibration uncertainty ($0^\circ.5$). We assigned the difference between this simulation and the fiducial angles as a systematic error. This error was the largest source of spurious B modes. We used two methods to estimate the error due to miscalibration of the relative angles. First, we ran a simulation using detector angles from Tau A instead of the fiducial model. Second, we randomly varied the detector angles by the relative calibration uncertainty (scatter among the diodes, $3^\circ.2$). These two methods gave consistent estimates for the systematic error. Finally, we considered possible time variation of the angles by simulating the effects of randomly varying the angles with $2^\circ.7$ RMS, based on the scatter of Tau A measurements. Detector-angle errors tend to decrease C_ℓ^{EE} and generate spurious C_ℓ^{EB} . The error for C_ℓ^{EB} can have either sign.

To estimate the systematic errors from pointing-model miscalibration, we compared different measurements of the pointing-model scatter. Using the same observations of Jupiter, RCW 38, and the Galactic center from which we created the pointing model, the RMS scatter was $2.17'$ (§5.1.4). If we instead used Moon observations to measure the scatter, the RMS was $1.6'$. We corrected for the pointing miscalibration at first order by including a pointing-smoothing term in the beam window function (§5.4.2). We assigned the systematic error corresponding to the difference between smoothings of $2.17'$ and $1.6'$. We added this uncertainty to the beam-shape systematic errors. Like the beam-shape miscalibration, the

pointing-model systematic error was a multiplicative effect for EE.

5.7.2 Instrumental Polarization

Instrumental imperfections created spurious polarization signals from unpolarized inputs. We did not correct for this effect in analysis. Instead, we showed that the resulting systematic error was small enough. First, we measured the $I \rightarrow Q/U$ leakage for each diode from the mini sky dips (§4.2.2 and [Tajima \(2011a\)](#)). The median leakage was 0.19%. The statistical uncertainty of each measurement was 0.29%. The systematic uncertainty was 0.16%, which was dominated by the dependence on the elevation at which we performed the sky dip (0.14%). We confirmed the mini-sky-dip measurements agree with measurements from large sky dips, Jupiter, and Tau A at the 0.03% level. Using the measured leakage, we simulated the corresponding spurious polarization in our analysis assuming CMB spectra from Λ CDM²⁹. We conservatively varied the leakage with elevation to account for the systematic error above. We also fluctuated the leakage based on the statistical measurement uncertainty. We assigned the spurious polarization in this simulation as a systematic error due to instrumental polarization. The combination of deck and sky rotation was essential in making the instrumental-polarization systematic error small enough.

The leakage considered above had the simplest beam shape, a monopole. We also considered leakage that generated spurious polarization signals with a dipole or quadrupole pattern ([Tajima, 2011b](#)). We measured each pattern from observations of Jupiter. The average leakage was 0.47% for the dipole pattern and 0.32% for the quadrupole pattern. As above, we simulated CMB temperature maps and convolved them with these leakage patterns to compute expected spurious polarization maps. We assigned systematic errors from the power spectra of these maps. The monopole leakage dominated at low ℓ ($\lesssim 300$); however, the

29. Although *WMAP* provides a high-signal-to-noise map of the CMB temperature, it does not measure the polarization well in our patches. Therefore, using *WMAP* as the input for this simulation would not provide a good estimate of the correlation between the real and spurious polarization.

dipole contribution was comparable at higher ℓ . The quadrupole contribution was negligible at all ℓ .

5.7.3 Type-B-glitch Miscorrection

Mis-correction of Type-B glitching caused additional instrumental polarization. I used the functional form of the Type-B correction to identify the possible effects. Then I used simulations to quantify the effects on the power spectra. I found that Type-B-glitch mis-correction increased the instrumental-polarization systematic error by 8% (Buder, 2011c).

If I did not use accurate Type-B correction parameters, some Type-B effects remain. I summarize the Type-B correction (§5.1.5) as

$$x'_{\text{DE}} = G(F(x_{\text{DE}}, x_{\text{TP}}; \vec{P}); \vec{P}'), \quad (5.54)$$

where x'_{DE} are the demodulated data after correction, x_{DE} are the demodulated data that would have been collected without Type-B glitching, x_{TP} are the TP data without Type-B glitching, \vec{P} is the set of glitch parameters, and \vec{P}' is the set of correction parameters. The collected data with Type-B glitching are $(y_{\text{DE}}, y_{\text{TP}}) = F(x_{\text{DE}}, x_{\text{TP}}; \vec{P})$. The function F models the effect of Type-B glitching; \vec{P} specifies the glitch locations and heights. The function G is the Type-B correction function. If $\vec{P} = \vec{P}'$ then $x' = x$; however, due to measurement uncertainty, $\vec{P}' \neq \vec{P}$. Expanding G near \vec{P} , the first order mis-correction causes

$$\frac{\partial x'_{\text{DE}}}{\partial x_{\text{DE}}} \neq 1 \quad (5.55)$$

$$\frac{\partial x'_{\text{DE}}}{\partial x_{\text{TP}}} \neq 0. \quad (5.56)$$

Eq. 5.55 represents the responsivity change introduced by mis-correction. Eq. 5.56 increases

the instrumental polarization because $x_{\text{TP}} \propto I$. Without Type-B, the $I \rightarrow Q/U$ leakage is

$$L_{\text{normal}} = \frac{dx_{\text{DE}}}{dx_{\text{TP}}}. \quad (5.57)$$

With Type-B, the total leakage is $\frac{dx'_{\text{DE}}}{dx_{\text{TP}}}$. The additional leakage due to Type-B is

$$\begin{aligned} L_{\text{Type-B}} &= \frac{dx'_{\text{DE}}}{dx_{\text{TP}}} - \frac{dx_{\text{DE}}}{dx_{\text{TP}}} = \frac{\partial x'_{\text{DE}}}{\partial x_{\text{DE}}} \frac{dx_{\text{DE}}}{dx_{\text{TP}}} + \frac{\partial x'_{\text{DE}}}{\partial x_{\text{TP}}} \frac{dx_{\text{TP}}}{dx_{\text{TP}}} - \frac{dx_{\text{DE}}}{dx_{\text{TP}}} \\ &= \frac{\partial x'_{\text{DE}}}{\partial x_{\text{TP}}} + \left(\frac{\partial x'_{\text{DE}}}{\partial x_{\text{DE}}} - 1 \right) \frac{dx_{\text{DE}}}{dx_{\text{TP}}}. \end{aligned} \quad (5.58)$$

The first term is the same as Eq. 5.56. The second term is a higher-order correction due to the responsivity change. Since (as I will show below) the responsivity change is $< 1\%$ and the normal and Type-B leakages are comparable, I ignored the correction and used Eq. 5.56 to estimate the Type-B leakage.

I ran simulations to calculate the Type-B leakage and responsivity change. I ran two sets of simulations: one with the fiducial correction parameters \vec{P} and another with fluctuated parameters. The difference between them is the Type-B miscorrection. I only simulated TOD, not the pipeline map making or power-spectra calculation. Furthermore, I used only a representative set of CES-diodes with the largest ‘‘Type-B badness’’: the absolute value of the difference of the Type-B χ^2 (§5.3.2) before and after correction. I tried many sets of fluctuated parameters including glitch locations, heights, and preamp factors differing from the measured ones by more than the measurement uncertainty. I also checked that the number of iterations (3) in the Newton’s method calculation of ΔG is sufficient. From the simulated TOD, I calculated the slope $\frac{\partial x'_{\text{DE}}}{\partial x_{\text{TP}}}$ (Figure 5.21). I saved the maximum absolute slope in each simulation as an upper limit on the possible leakage. I also calculated $\frac{\partial x'_{\text{DE}}}{\partial x_{\text{DE}}}$ and saved the maximum slope change, including its sign³⁰. Finally, I calculated the Type-B

30. The absolute responsivity systematic error depends on whether the sign is the same for different CES-

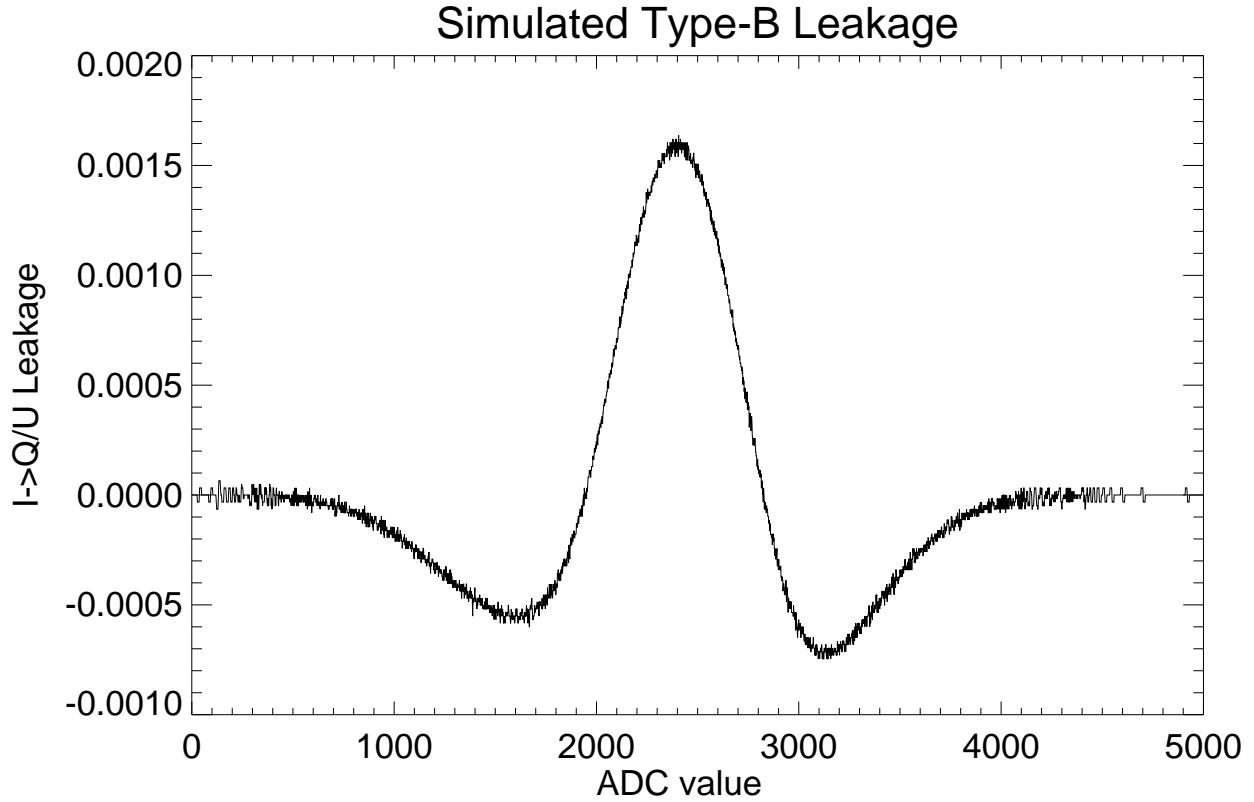


Figure 5.21: Type-B-glitch miscorrection created instrumental polarization that varies with the TP offset. The units of the horizontal axis are arbitrary. The average polarization is zero; however, the fluctuations create spurious polarization power.

χ^2 from the simulated TOD.

I used the relationship between the Type-B χ^2 and simulated leakage to propagate the leakage to the power spectra. I computed an effective Type-B leakage for each CES-diode from its Type-B χ^2 . Although there are many other reasons for excess χ^2 , attributing all of it to miscorrection leads to the most conservative estimate. I computed the effective leakage due to Type-B miscorrection

$$L_{\text{eff}}^{\text{Type-B}} = \sum_{\text{CES-diodes}} L_i / N, \quad (5.59)$$

diodes.

where N is the total number of CES-diodes and L_i is the effective leakage for each CES-diode given by:

1. If the Type-B badness was less than 0.1, I set the effective leakage to zero. This excluded CES-diodes that did not have significant Type-B. The result did not depend strongly on the threshold (Figure 5.22).
2. If the χ^2 of this CES-diode was less than 1, I set the effective leakage to zero. These CES-diodes were well corrected, and the nature of the χ^2 simulation forced the simulated χ^2 to be greater than 1.
3. If the χ^2 of this CES-diode was greater than the maximum χ^2 in simulation, the maximum simulated leakage was the effective leakage. Since I simulated extreme parameter fluctuations, if the χ^2 of this CES-diode was larger than any of them, it cannot be the result of miscorrection. I adopted the maximum leakage as a conservative limit on the miscorrection that may be hidden by whatever other effect caused the bad χ^2 .
4. If none of the above apply, I interpolated the simulated leakage- χ^2 relationship to find the effective leakage at this χ^2 .
5. I repeated the steps above for each type of parameter fluctuation and used the largest resulting leakage. This was the conservative choice because the data χ^2 did not indicate which Type-B parameter was inaccurate.

The resulting $L_{\text{eff}}^{\text{Type-B}} = 0.0012$. In the power spectrum, this leakage effect added in quadrature with the normal leakage. This is because the average Type-B leakage was zero. The leakage due to miscorrection is a sum of odd functions³¹ so the average leakage must be zero as long as x_{TP} was sampled uniformly. Since the fluctuation of x_{TP} during the season was larger than the glitch width and there is no reason to prefer one side of the glitch, there

31. See Eq. 59 of [Kusaka \(2010c\)](#).

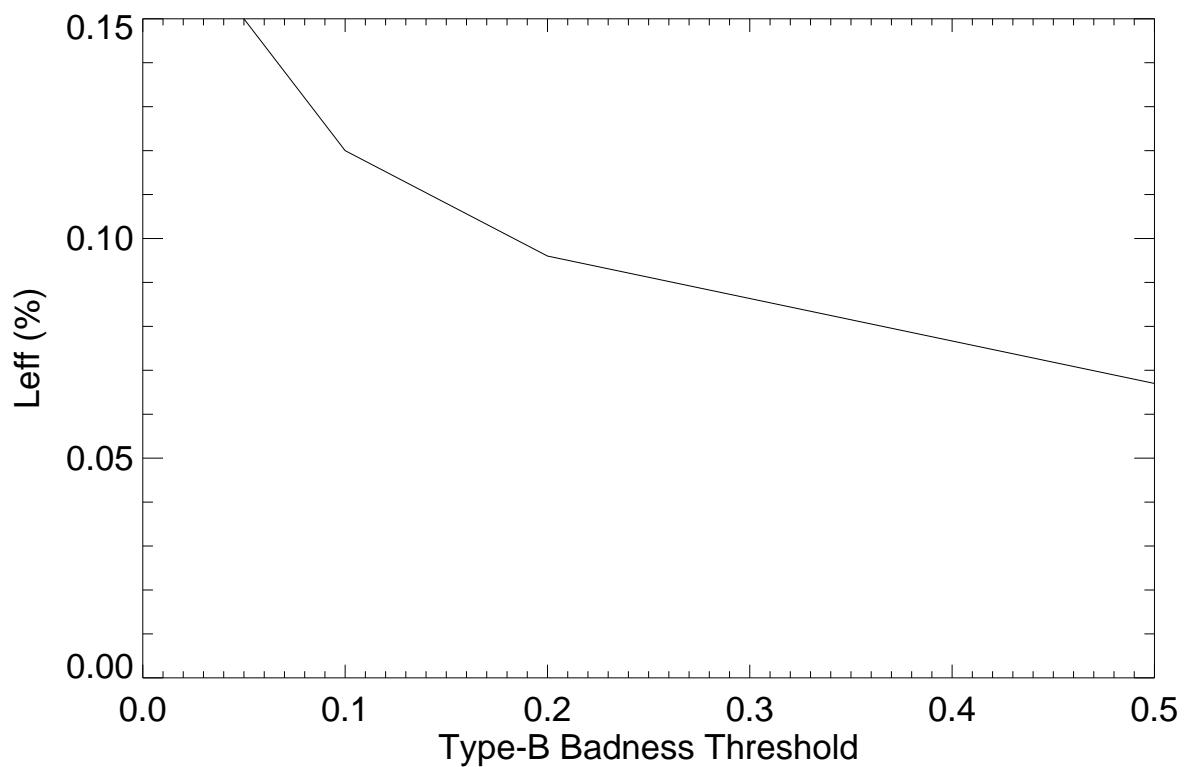


Figure 5.22: The systematic error from Type-B miscorrection did not depend strongly on the threshold I used to exclude CES-diodes without significant Type-B glitching.

was effectively uniform sampling. The average miscorrection leakage was zero, so there was no correlation with the normal leakage. The main effect was the variance produced by the correlation of the miscorrection leakage with itself. I estimated this effect by comparing the effective leakage to the normal leakage (0.19%). Since the effect in the power spectra was proportional to the leakage squared, the effect of miscorrection was $(0.12/0.19)^2 \approx 40\%$ of the normal leakage effect. Since the two effects added in quadrature, the total systematic error was $\sqrt{1 + 0.4^2} \approx 1.08$ times that from normal leakage alone.

The other effect of Type-B miscorrection was the responsivity change; however, I found that the systematic error was negligible. I computed an effective responsivity change

$$\Delta R_{\text{eff}}^{\text{Type-B}} = \sum_{\text{CES-diodes}} \Delta R_i / N \quad (5.60)$$

using the same method as for leakage. ΔR is the fractional responsivity difference caused by Type-B miscorrection. I also computed a relative-responsivity fluctuation

$$\sigma R_{\text{eff}}^{\text{Type-B}} = \sqrt{\sum_{\text{CES-diodes}} \Delta R_i^2 / N}, \quad (5.61)$$

which quantifies the possible relative-responsivity uncertainty introduced in the average CES-diode. The results were $\Delta R_{\text{eff}}^{\text{Type-B}} = -0.002$ and $\sigma R_{\text{eff}}^{\text{Type-B}} = 0.005$. The former created a negligible absolute-responsivity error (0.2%). The relative-responsivity error of 0.5% was small compared to the scatter in the relative-responsivity model (12.5%).

As mentioned in §5.1.5, the Type-B correction depended on the noise being Gaussian. To test this assumption, I stacked 800-kHz snapshot data for one diode (Buder, 2011a). To differentiate between inherent non-Gaussianity and non-Gaussianity introduced by Type-B miscorrection, I chose a diode without Type-B glitching. I accumulated one month of snapshots (284071 samples). To normalize the data, I subtracted the mean from each snapshot and divided by the standard deviation. Although I detected small skewness and excess

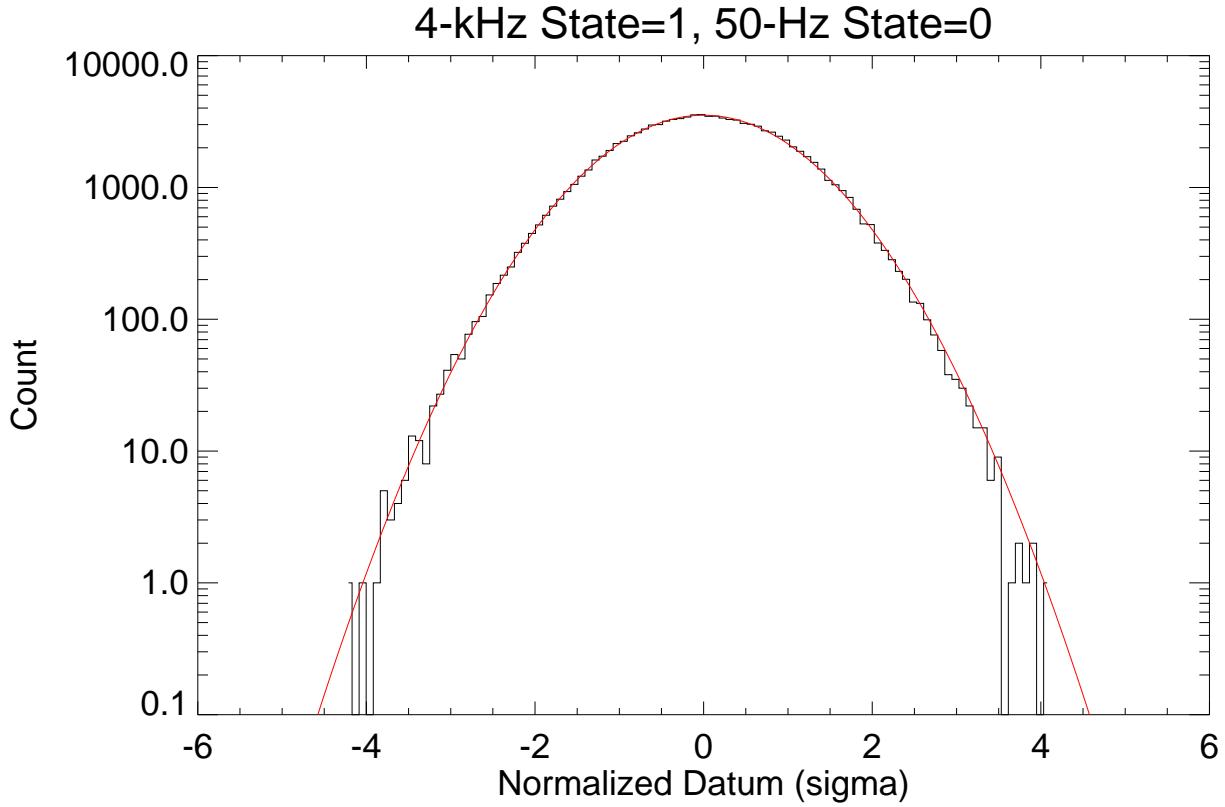


Figure 5.23: The 800-kHz data were nearly Gaussian-distributed. If the data had significant non-Gaussianity, I would have had to account for it in Type-B-glitch correction. The red line is a Gaussian distribution with mean 0 and RMS 1.

kurtosis, the normalized data were well-modeled by a Gaussian distribution (Figure 5.23).

5.7.4 Sun Pickup

Optics sidelobes seeing the Sun could create systematic errors. We found contamination from the Sun when it passed through known sidelobes before we installed the upper ground screen, which eliminated the contamination. Although we rejected data with such contamination, the rejection may not be perfect and could have an associated systematic error. To estimate it, we first made a map of CMB data in Sun-centered coordinates (QUIET Collaboration, 2012b; Chinone, 2011b). This map enhanced the Sun signal compared to the normal analysis and used the Sun as a sidelobe-finding source. The strongest signal in this map was < 1 mK.

After data selection, we did not find any signal in the Sun-centered map. To estimate the possible systematic error, we performed simulations including Sun contamination drawn from the Sun-centered map. In these simulations we did not reject the contaminated data, so the error estimate was conservative. We assigned the power-spectra differences between simulations with and without Sun contamination as the possible systematic errors. We also considered the effect of the Moon. However, the Moon was an order of magnitude fainter than the Sun, and the Sun contamination was small. We concluded that possible Moon-pickup systematic errors were negligible.

5.7.5 Type-O Glitching

We found a type of polarized scan-synchronous signal (“Type-O glitching”) whose residual in analysis created a systematic error. We first discovered this effect as an elevation-dependent polarization offset in large-sky-dip data (Tajima, 2011a). In CMB data we found a corresponding polarized azimuth structure in some CES-diodes. The signal depended strongly on the azimuth and deck of the observation (Nagai, 2012a). In filtering, we subtracted an arbitrary function of azimuth from each CES-diode’s data (§5.3.1). This filter removed any scan-synchronous signal whose amplitude and azimuth-dependence were constant for the duration of a CES. Therefore, only scan-synchronous signals whose amplitude or azimuth-dependence changed during a CES could contaminate the result. We did not detect any variation of the Type-O signal within a CES; however, because of statistical uncertainty we could only place an upper limit on the possible variation. We ran simulations with the scan-synchronous signal varying within that limit and assigned the resulting spurious polarization as a systematic error (Nagai, 2012d).

5.7.6 Negligible Effects

In addition to the assigned systematic errors above, we considered many other effects that we determined were negligible. Here I discuss some examples. We considered effects from ground signals that remained after filtering. We considered the impact of compression, the change of the responsivity as a function of input power³². Finally, we considered whether foregrounds contributed significantly to our power spectra.

To limit the possibility of significant ground signal, we removed a slope from each half-scan and an azimuth-dependent signal from each CES-diode's data. Since we limited the azimuth dependence to 40 bins in azimuth, high-spatial-frequency components of a ground signal might remain after the filter. Alternatively, a ground signal that varied in time and had more than linear structure would remain after the filter. I discuss these two possibilities below.

To address the finite binning of the ground filter, I re-ran the analysis with 80 filter bins. This analysis was consistent with the fiducial one. The differences were $< 2\%$ of the statistical uncertainty, on average (Buder, 2012b).

To consider the possible time variation of the ground signal, we ran a simulation based on the CES-diode with the worst ground pickup. Although we did not detect any time variation in the data, we created a varying ground signal using that CES-diode as a template and allowing the signal to vary within its statistical uncertainty. We added that signal to the simulation of one CES and computed the resulting power spectra. We scaled the simulation to the full season based on the number of CESes and the amount of ground signal in each. The result was negligible compared to the other systematic errors (Nagai, 2012c).

Emission from diffuse Galactic sources or compact extra-Galactic sources could have contaminated the results with non-CMB polarization power. We limited the diffuse synchrotron

32. See Appendix 10.1 of QUIET Collaboration (2012b).

foregrounds in our patches to the level of $r = 0.05$ from our previous Q-band measurement (QUIET Collaboration et al., 2011). Lower-frequency *WMAP* observations supported this limit. It was well below our statistical uncertainty. We used the *Planck Sky Model* (Delabrouille et al., 2012) to predict the thermal dust emission in our patches. The prediction was below our statistical uncertainty (Tajima, 2012). We also cross-correlated the model with the QUIET maps and did not detect significant power in the cross-correlation.

We masked the two most polarized compact sources in our maps (§5.4.2). To consider the effect of other known sources, I re-ran the analysis masking other sources known from the *WMAP* (Gold et al., 2011) and *Planck* (Planck Collaboration, 2011) catalogs. I masked all pixels within $0^{\circ}5$ of a *WMAP* source. I masked all pixels within $0^{\circ}5$ of a source from the *Planck* 100-GHz catalog; I increased the mask radius (by twice the catalog major axis FWHM) if *Planck* found the source was extended. I did not find a significant difference between this result and the fiducial analysis. The largest difference was a C_{ℓ}^{EB} shift corresponding to 17% of the statistical error in patch 7b, the least observed patch (Buder, 2012a). Moreover, the statistical significance of this shift was only 2.5σ . Even if this effect were real, it would only increase the total error by $< 1\%$. This test confirmed that we did not need to mask the next-brightest class of sources; however, the sensitivity of *WMAP* and *Planck* limited this test. The flux density limit of the *Planck* catalog was 344 mJy, approximately $90\mu\text{K}$ (antenna temperature) for QUIET. Assuming a conservative polarization fraction of 10%, *Planck* would find all sources down to a polarization of $9\mu\text{K}$. We used models of polarized compact sources (Battye et al., 2011; Tucci & Toffolatti, 2012) to estimate the contribution of sources below this detection threshold. It was below the level of $r = 0.01$.

Chapter 6

Results and Discussion

Scientific discovery and scientific knowledge have been achieved only by those who have gone in pursuit of it without any practical purpose whatsoever in view.

Max Planck

After finalizing all analysis choices in the blind stage of analysis, we revealed the CMB power spectra. We reported these results as frequentist confidence intervals. We confirmed they were consistent with the results of the alternate QUIET analysis. We then evaluated the agreement of our results with the standard Λ CDM cosmological model. Towards the primary science goal of constraining inflation, we set a limit on the tensor-to-scalar ratio $r = 1.48^{+1.06}_{-0.92}$. Although this limit was not competitive with other results including temperature anisotropies, it had the lowest systematic error of any limit using CMB polarization information alone.

6.1 Confidence Intervals

Using fiducial simulations of varying CMB power levels, we constructed a likelihood function, \mathcal{L} , following the method of [Hamimeche & Lewis \(2008\)](#).

$$-2 \ln \mathcal{L}(C_b | \hat{C}_b) \equiv \sum_{bb'} [g_b(\hat{x}_b) C_{fb}] [M_f^{-1}]_{bb'} [C_{fb'} g_{b'}(\hat{x}_{b'})], \quad (6.1)$$

where

$$\hat{x}_b \equiv \frac{\hat{C}_b + N_b^{\text{eff}}}{C_b + N_b^{\text{eff}}} \quad (6.2)$$

$$g_b(x) \equiv \text{sign}(x - x_m) \sqrt{2\sigma_b^2 \ln \left(\frac{x_m S_b(x_m)}{x S_b(x)} \right)} \quad (6.3)$$

$$S_b(x) \equiv P_{\chi^2}^M(x - 1; \nu_b, \sigma_b) \quad (6.4)$$

$$\sigma_b = \frac{\sqrt{[M_f]_{bb}}}{C_{f_b} + N_b^{\text{eff}}} \quad (6.5)$$

$$P_{\chi^2}^M(x; \nu, \sigma) \equiv \frac{\sqrt{2\nu}}{\sigma} P_{\chi^2} \left(\nu(x\sqrt{2\nu}/\sigma + 1); \nu \right) \quad (6.6)$$

C_{f_b} are the bandpowers from fiducial simulations, M_f is the covariance matrix of \hat{C}_b from fiducial simulations, x_m is the value of x that maximizes $xS(x)$, $P_{\chi^2}(x; \nu)$ is the probability density function for a χ^2 distribution with ν degrees of freedom, N_b^{eff} is an instrument-noise term from simulations, and ν_b is an effective number of degrees of freedom we fit from simulations. (Eq. 6.1 is sufficient for EE or BB. For the EB likelihood we used [Hamimeche & Lewis \(2008\)](#) Eq. 47.) We calibrated \mathcal{L} using simulations without any CMB power, with power equal to Λ CDM, and with non-zero B-mode power ([PCL group, 2010](#)). The limited number of simulations introduced an uncertainty of 10% on the statistical error. We combined the four patches by multiplying their likelihood functions.

From this function we constructed frequentist two-sided 68% confidence intervals (Figure 6.1). We used the C_b that maximize the likelihood as the central values. We chose the confidence intervals such that their boundaries are the points where $-2 \ln \mathcal{L}$ increases by 1 from its minimum. When computing the confidence interval for any one b , we allowed all $b' \neq b$ to vary to maximize \mathcal{L} . When computing patch-combined intervals for one spectrum (e.g. EE) we did not force the other spectra (e.g. BB and EB) to be the same in all patches. The goodness of fit for the patch-combined central values was $\chi_{\min}^2 \equiv -2 \ln \mathcal{L}_{\min} = 48.3$

(EE), 77.0 (BB), and 64.4 (EB) for $19 \times 3 = 57$ degrees of freedom. The corresponding PTEs were 79%, 4%, and 23%, showing that the individual patch results were consistent with each other (Figure 6.2). Because of the limited sky coverage, neighboring ℓ bins had ≈ -0.1 correlation. Bins that were further apart were effectively uncorrelated. There were also non-negligible correlations among EE, BB, and EB in the same bin. The “window function” $K_{bb'}^{-1} P_{b'\ell} M_{\ell\ell'} B_{\ell'}^2$ quantifies the dependence on C_ℓ of each bandpower (Figure 6.3).

6.2 Comparison with Alternate Analysis Pipeline

As a final check before interpreting the results, I compared the results of the two parallel QUIET analyses (Figure 6.4). I computed the χ^2 for the hypothesis that the two analyses were measurements of the same spectra with independent statistical fluctuations. The total χ^2 was 8.4, 6.5, and 9.1 for EE, BB, and EB, respectively, each with 19 degrees of freedom. The corresponding PTEs were 98%, 100%, and 97% indicating that our two analyses of the same data were much more consistent than independent analyses of disjoint data sets. In this comparison I adjusted the results of the alternate analysis to have the same absolute responsivity¹. Without making that adjustment, the PTEs were 89%, 100%, and 97%. Furthermore, both pipelines found consistent results for all parameter estimates.

6.3 Consistency with Λ CDM

After confirming that our analysis was internally consistent, we compared our results to the standard cosmological model (Λ CDM with parameters from Komatsu et al. (2011)). The χ^2 to Λ CDM was 24.3, 21.5, and 21.5 for EE, BB, and EB, respectively. There were 19 degrees of freedom for each spectrum. The corresponding PTEs were 19%, 31%, and 31%. Because the responsivity systematic error was significant, we allowed the absolute

1. In the blind stage of analysis, we found that the pipelines’ absolute-responsivity calibration differed by 3.5% (7% in C_ℓ).

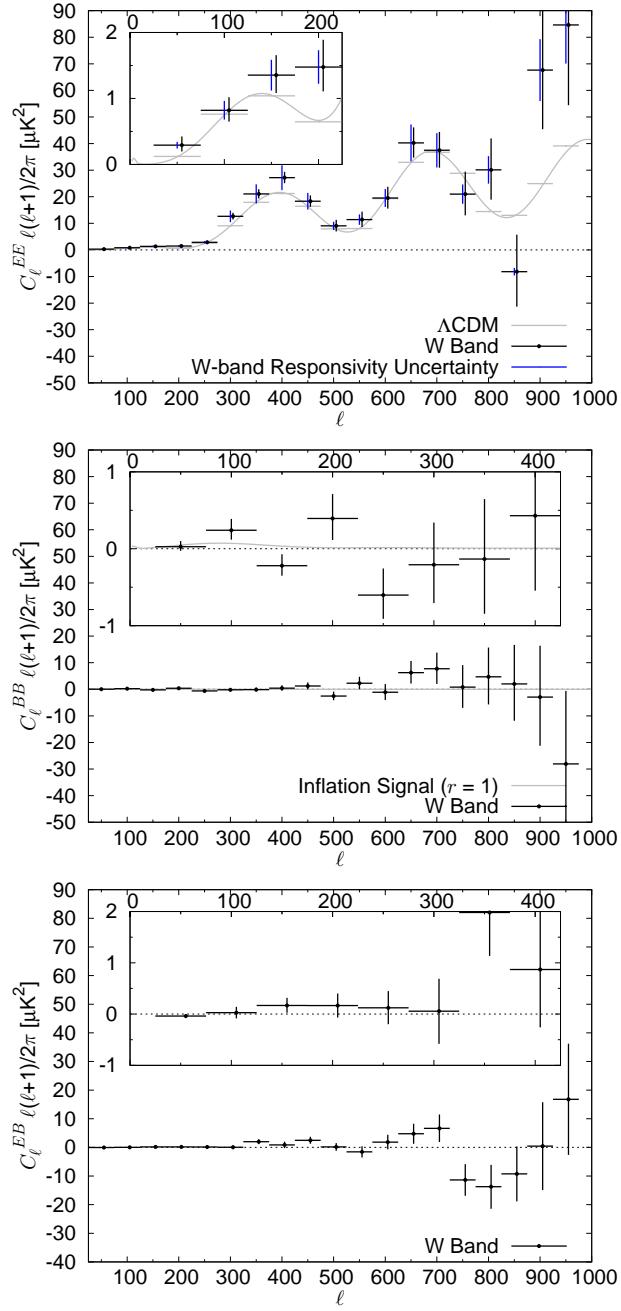


Figure 6.1: We reported the measurements of the CMB power spectra as frequentist two-sided 68% confidence intervals. Insets show the low- ℓ region. For EE, the blue error bars show the absolute-responsivity uncertainty, which dominates at some ℓ . The horizontal displacement of points is for visualization purposes only. The horizontal gray lines show the expected C_b^{EE} in the standard model (Λ CDM).

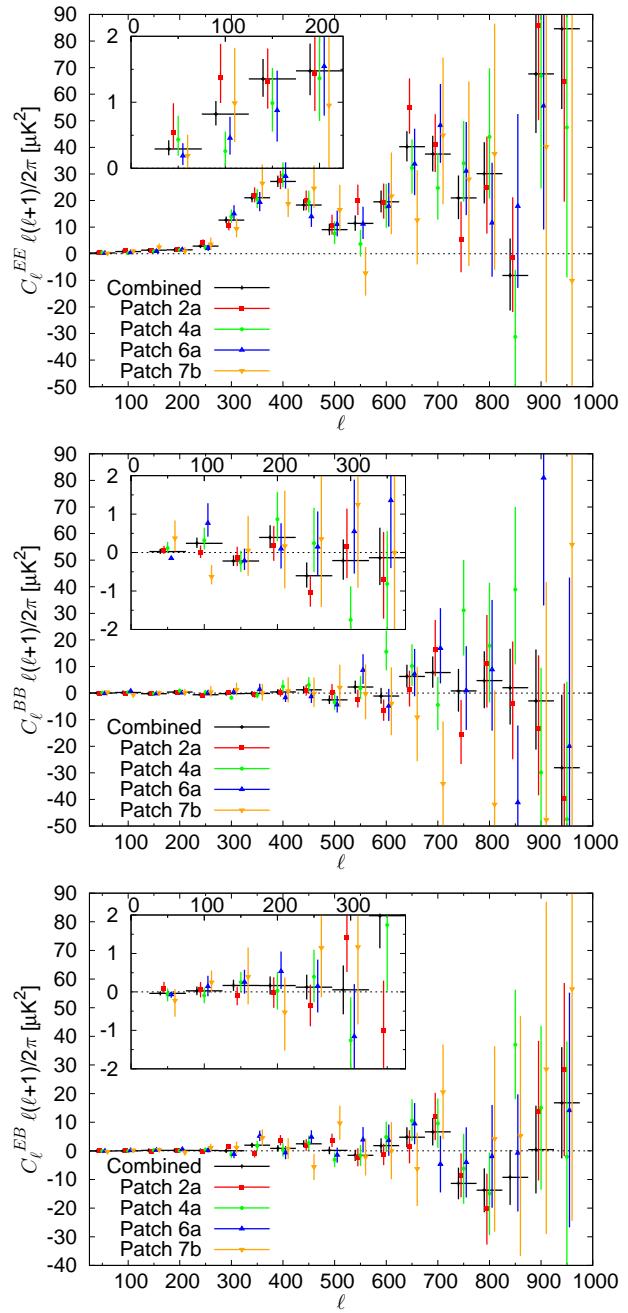


Figure 6.2: The spectra from each patch were consistent with the combined spectra. The horizontal displacement of points is for visualization purposes only.

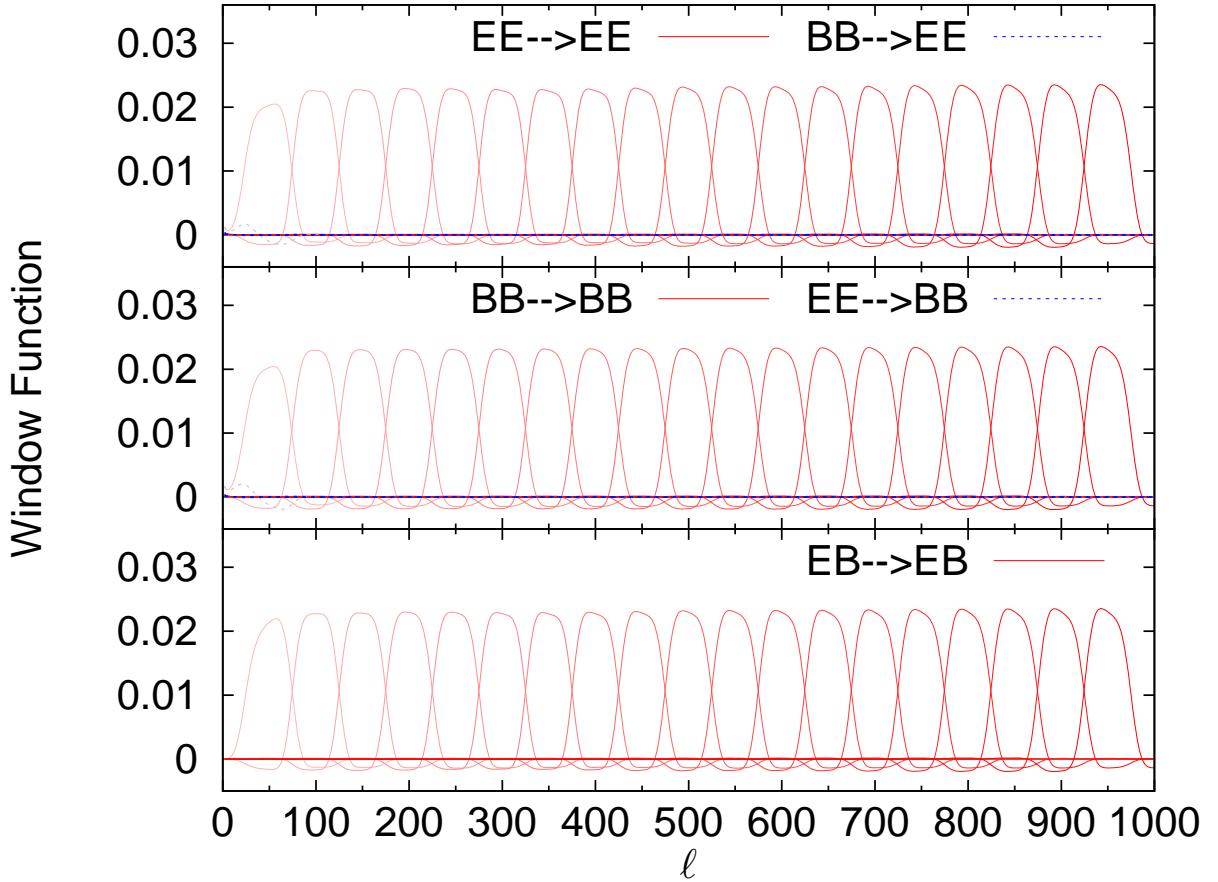


Figure 6.3: The window functions show the ℓ weighting of each bandpower, including the effects of the beam and mode coupling. Each bandpower had some response in the neighboring band, resulting in 10% correlation between neighboring ℓ bins. The response of our C_ℓ^{EE} estimator to C_ℓ^{BB} was small, and vice versa. I averaged the window functions for all four patches.

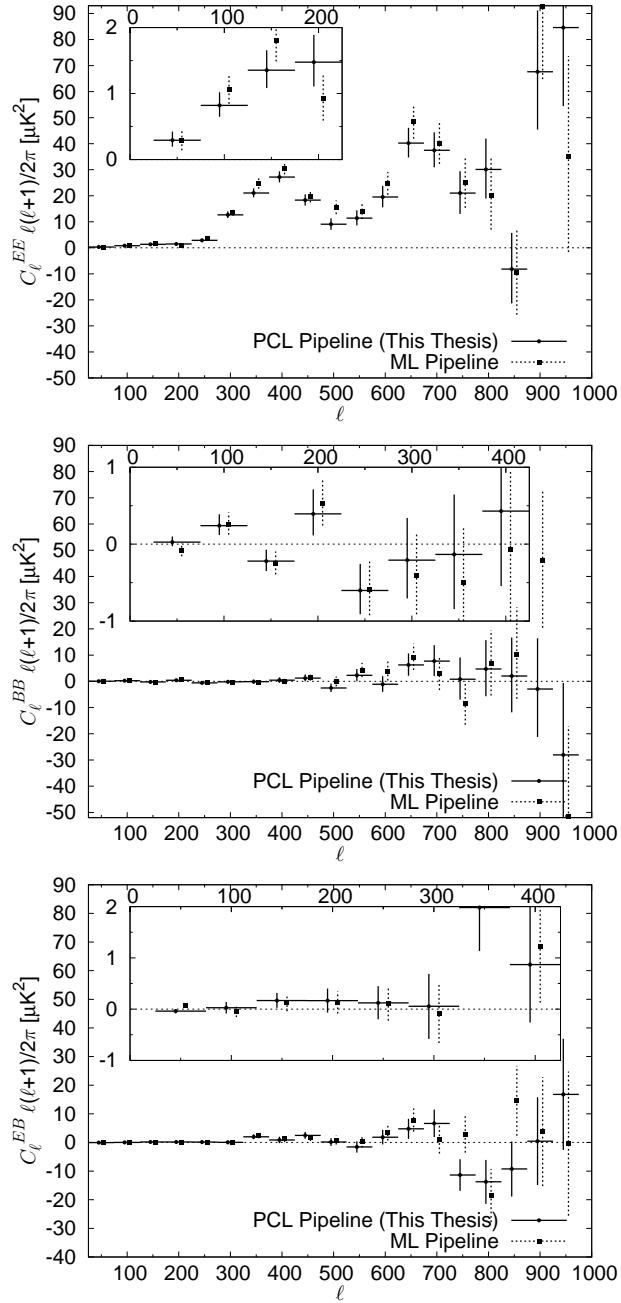


Figure 6.4: We computed the CMB power spectra with two parallel analyses; they found consistent results. The alternate analysis (ML Pipeline) used a different responsivity model, resulting in 7% higher absolute responsivity for the power spectra. The horizontal displacement of points is for visualization purposes only.

responsivity to vary when calculating this consistency. We used our fiducial responsivity as a constraint, assuming it had a Gaussian distribution with RMS equal to the systematic error we assigned. If we instead allowed the responsivity to be a free parameter, the PTE for EE was 20%. Similarly, we included the dominant systematic error due to detector-angle miscalibration when calculating the EB consistency. We found no evidence for deviation from Λ CDM or for non-zero BB or EB power.

6.4 Limit on Inflation

We constrained inflationary models using only the B-mode measurements at $25 \leq \ell \leq 175$ (the first three bandpowers). We fit a B-mode template spectrum with variable r to these bandpowers. For simplicity, we fixed $n_t = 0$ in the fit; that choice made the template directly proportional to r . We found $r = 1.48^{+1.06}_{-0.92}$ as a 68% confidence interval, corresponding to $r < 3.3$ as a 95% confidence upper limit. We calculated the upper limit to be the point where $-2 \ln \mathcal{L}$ increased by 2.69 from its minimum. Statistically, this measurement did not improve on the existing limit $r < 0.2$ (Komatsu et al., 2011). However, our limit used CMB B-mode polarization alone with smaller systematic error than ever before.

6.5 Conclusions

We measured the CMB polarization with smaller errors than ever before. Our C_ℓ^{EE} measurement was the most sensitive so far between $300 \lesssim \ell \lesssim 500$ (Figure 6.5). We placed new limits on C_ℓ^{BB} and C_ℓ^{EB} . The most important result from QUIET was a new, low-systematic-error measurement of the tensor-to-scalar ratio r from C_ℓ^{BB} alone. The statistical power of this measurement was insufficient to improve on existing constraints. However, using a combination of techniques (especially double demodulation, a side-fed Dragonian telescope, an absorbing ground screen, deck rotation, cross-correlation, and blind analysis) we reduced

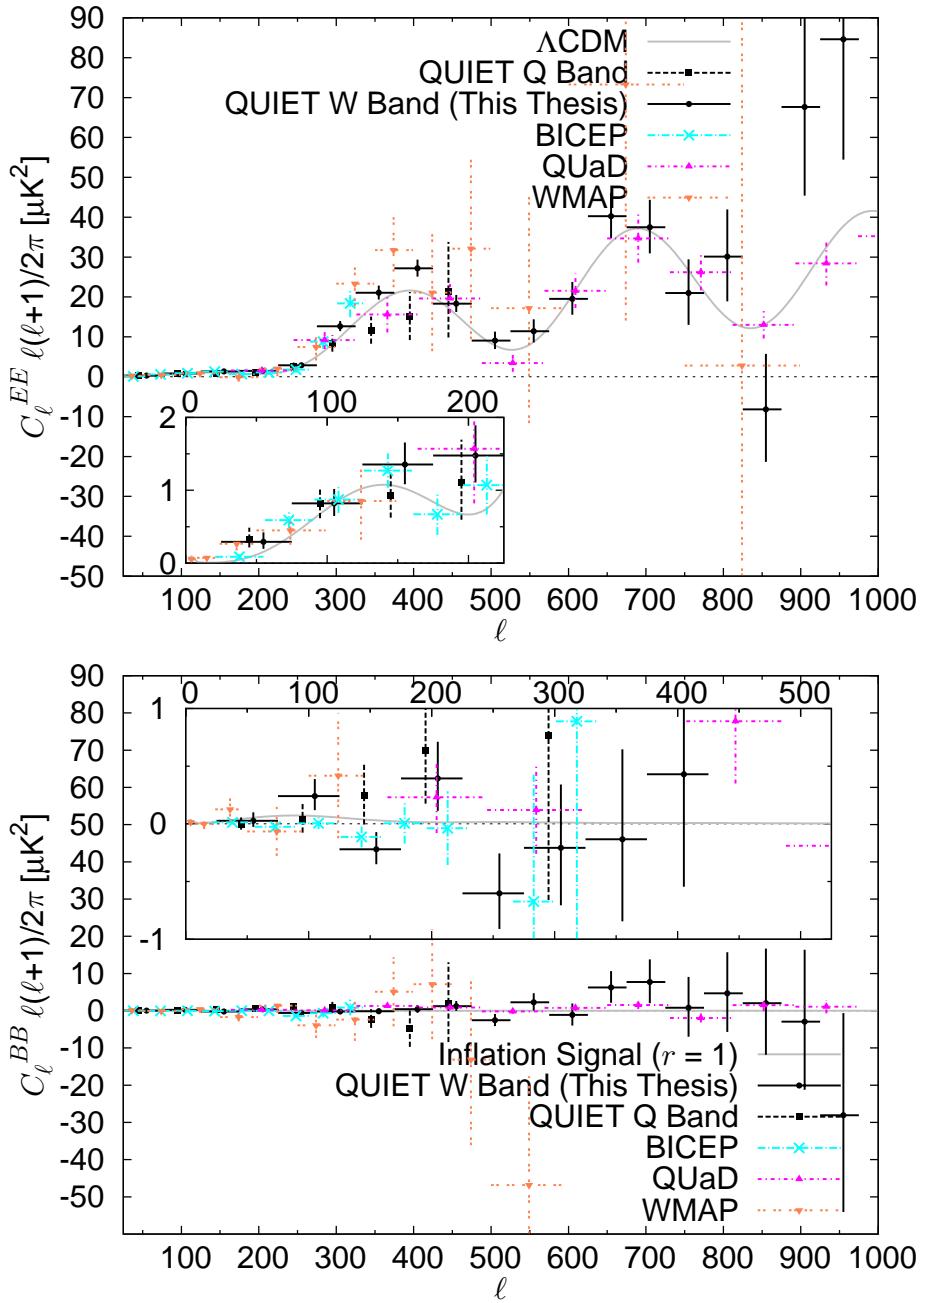


Figure 6.5: We made the most sensitive measurement of the CMB polarization between $300 \lesssim \ell \lesssim 500$. I only show statistical errors. QUIET Q and W bands share horizontal error bars; the horizontal displacement of points is for visualization purposes only.

systematic errors by more than an order of magnitude to below the level of $r = 0.01$. Already, new projects are employing these techniques with even more sensitive polarimeter arrays to search for new evidence of inflation in the CMB.

Appendix A

DataCompilation Status Flags

This appendix lists the meaning of the Online Software (`DataCompilation`) status flags.

- `DATA_STATUS_UNEVEN_DEMOD` means that the number of samples was not the same in the two 4-kHz phase-switch states. Such data are useless.
- `DATA_STATUS_NO_ADC_CONN` means that `DataCompilation` cannot connect to `adc_server`. Data will not be collected until connectivity is restored.
- `DATA_STATUS_NO_PER_CONN` means that `DataCompilation` cannot connect to `PeripheralServer`. Some housekeeping data will be missing until connectivity is restored.
- `DATA_STATUS_NO_BIAS_CONN` means that `DataCompilation` cannot connect to `BiasServer`. Biasing data will be missing until connectivity is restored.
- `DATA_STATUS_NO_AGG_CONN` meant that `DataCompilation` cannot connect to `DataArc`. Under most circumstances this is of no concern because we only used `DataArc` in special circumstances.
- `DATA_STATUS_NO_DB_CONN` means that `DataCompilation` could not read the Receiver Database. If this flag remains on it may indicate the database is corrupt, missing, etc.
- `DATA_STATUS_BUFFER_FULL` means that the `DataCompilation` data buffer became full and data were discarded.
- `DATA_STATUS_INT_ERR` indicates a critical error in `DataCompilation` requiring a restart of the Online Software.

- `DATA_STATUS_UNEVEN_MASK` was supposed to indicate that the blanking mask was not the same on all ADC Boards; however, due to firmware limitations there are too many false alarms for this flag to be useful.
- `DATA_STATUS_UNEVEN_CLOCKS` was supposed to indicate that clocks were not the same on all ADC Boards; however, due to firmware limitations there are too many false alarms for this flag to be useful.
- `DATA_STATUS_BAD_SCALE` indicates the scale factor (number of 800-kHz samples accumulated in each 100-Hz sample) is impossible. Usually this means the firmware is not working.
- `DATA_STATUS_HK_PARSE` means `DataCompilation` could not parse the housekeeping data from `adc_server`. Data were lost.
- `DATA_STATUS_DEMOD_PARSE` means that `DataCompilation` could not parse the radiometer data from `adc_server`. Data were lost.
- `DATA_STATUS_SNAP_PARSE` means that `DataCompilation` could not parse snapshot data from `adc_server`. At least 1 snapshot was lost.
- `DATA_STATUS_SERVER_THREAD` indicates an unrecoverable error in the `DataCompilation data_server` thread. The Online Software should be restarted.
- `DATA_STATUS_HK_THREAD` indicates an unrecoverable error in the `DataCompilation hk` thread. The Online Software should be restarted.
- `DATA_STATUS_SERVE_THREAD` indicates an unrecoverable error in the `DataCompilation data_serve` thread. The Online Software should be restarted.
- `DATA_STATUS_DEMOD_THREAD` indicates an unrecoverable error in the `DataCompilation demod` thread. The Online Software should be restarted.

- DATA_STATUS_INVALID_WPID_VALUE indicates an invalid value received from the WPID.
- DATA_STATUS_INVALID_CPID_VALUE indicates an invalid value received from the CPID.
- DATA_STATUS_INVALID_PRESSURE_VALUE indicates invalid data from the cryostat pressure sensor.
- DATA_STATUS_INVALID_TEMP_VALUE was supposed to indicate invalid data from the thermometers on the ground screen. However, there were too many false alarms.
- DATA_STATUS_NO_WPID_CONN indicated there was no connection to the WPID. Usually, the Electronics Enclosure temperature regulation also failed.
- DATA_STATUS_NO_CPID_CONN indicated there was no connection to the CPID. Usually, the cryostat temperature regulation also failed.
- DATA_STATUS_NO_PRESSURE_CONN indicated there was no connection to the cryostat pressure sensor. Data were missing.
- DATA_STATUS_NO_TEMP_CONN was supposed to indicate connection failure to the ground-screen thermometers. However, there were too many false alarms.
- DATA_STATUS_CRATE_PROBLEM indicated a problem with the ADC Backplane.
- DATA_STATUS_WPID_ALARM_1 was linked to the WPID “alarm” function, which we did not use.
- DATA_STATUS_WPID_ALARM_2 was linked to the WPID “alarm” function, which we did not use.

- `DATA_STATUS_``AGG_THREAD` indicated an unrecoverable error in the `DataCompilation archive_aggregator` thread. The Online Software should be restarted.
- `ADC_STATUS_``EXT_CLOCK_BAD` indicated that the Master ADC could not synchronize with the 10-MHz clock. Data taken in this situation were not properly synchronized.
- `ADC_STATUS_``NO_EXT_1HZ` indicated that the Master ADC was not using the 1-Hz clock for timing synchronization. Data taken in this situation were not properly synchronized.
- `ADC_STATUS_``NOT_TFP` indicated that `adc_server` could not provide timestamps from the Time-code Reader. Such data do not have reliable timing synchronization.
- `ADC_STATUS_``4HZ_COUNT_MATCH` indicated that the 4-Hz frame counter was not the same for all ADC Boards. Either firmware or `adc_server` problems could cause this flag.
- `TFP_STATUS_``NOLOCK` indicated that the Time-code Reader could not synchronize to the IRIG-B time code. The data were not properly synchronized.
- `TFP_STATUS_``BATTFAIL` indicated that the Time-code Reader battery failed.
- `TFP_STATUS_``TIMEOFF` indicated that the Time-code Reader reported a time offset. This created too many false alarms to be useful.
- `TFP_STATUS_``FREQOFF` indicated that the Time-code Reader reported a frequency offset. This created too many false alarms to be useful.
- `BSERV_STATUS_``BAD_COMMAND` indicated that `BiasServer` could not send a command because it was invalid.

- BSERV_STATUS_NO_ADC_CONN indicated that **BiasServer** could not connect to **adc_server**.
- BSERV_STATUS_BAD_ADC_COMM indicated a communication failure between **BiasServer** and **adc_server**. The most recent bias command failed.
- BSERV_STATUS_NO_ADC_REPLY indicated that **adc_server** never indicated success for a ‘b’ request (Appendix [B.2](#)). The bias command may have failed.
- BSERV_STATUS_ADC_MINOR_ERR indicated strange communication with **adc_server** that usually does not cause bias commands to fail.
- BSERV_STATUS_NO_BIAS_FRAME indicated that **BiasServer** did not have memory to store the bias data.
- BSERV_STATUS_MUTEX_ERROR indicated that **BiasServer** needed to be restarted.
- BSERV_STATUS_NO_RC_CONN indicated that **BiasServer** could not connect to **ReceiverControl**.
- BSERV_STATUS_BAD_RC_COMM indicated communication failure between **BiasServer** and **ReceiverControl**.
- BSERV_STATUS_BAD_DC_COMM indicated communication failure between **BiasServer** and **DataCompilation**.
- BSERV_STATUS_DATA_THREAD indicated an unrecoverable error in **BiasServer** requiring restart.
- BSERV_STATUS_COMMAND_THREAD indicated an unrecoverable error in **BiasServer** requiring restart.

- PERSERV_STATUS_COMMAND_THREAD indicated an unrecoverable error in `PeripheralServer` requiring restart.
- PERSERV_STATUS_BAD_DC_COMM indicated nothing.
- PERSERV_STATUS_DATA_THREAD indicated an unrecoverable error in `PeripheralServer` requiring restart.
- PERSERV_STATUS_MUTEX_ERROR indicated an unrecoverable error in `PeripheralServer` requiring restart.
- PERSERV_STATUS_NO_PER_FRAME indicated an unrecoverable error in `PeripheralServer` requiring restart.
- DATA_STATUS_DARC_THREAD indicated an unrecoverable error in `DataCompilation data_arc_service` thread requiring restart.
- DATA_STATUS_DARC_BUFFER_FULL indicated that the buffer of frames to send to `DataArc` became full. During normal operation this flag could be ignored.
- DATA_STATUS_CBI_BUFFER_FULL indicated that the buffer of frames to send to `cbicontrol` became full. Data were lost.
- DATA_STATUS_CBI_THREAD indicated an unrecoverable error in `DataCompilation cbi_service` thread requiring restart.
- DATA_STATUS_SS_ERROR indicated `DataCompilation` encountered an error sending data to `cbicontrol`.
- DATA_STATUS_NULL_FRAME indicated that `DataCompilation` inserted a null frame for timing alignment.
- DATA_STATUS_BAD_TIMING indicated that the data were not properly aligned.

- `DATA_STATUS_BAD_QUAD_SCALE` indicated that the scale factor for quadrature data was impossible.
- `DATA_STATUS_WPID_THREAD` indicated an error in `DataCompilation wpid_output` thread. The WPID output monitoring data were not recorded properly.
- `DATA_STATUS_NO_PDU` indicated that `DataCompilation` could not receive data from the Electronics Enclosure smart power strip.
- `DATA_STATUS_NOISE_SOURCE_ON` indicated that the (calibration) noise source was on.
- `BSERV_STATUS_BAD_PS_STATE` indicated that `adc_server` sent the wrong phase-switch state information to `BiasServer`. Resetting the phase-switch state usually fixed this problem.

Appendix B

Online Software Details

B.1 adc_server Packet Status Bytes

This section lists the contents of the 16 bytes of status information `adc_server` sent to `DataCompilation` in every radiometer and housekeeping packet.

1. The frame counter, a 4-byte integer.
2. A flag (1 byte) indicating whether the Time-code Reader battery failed.
3. A flag indicating whether the Time-code Reader is locked to the time code.
4. A flag indicating whether the Time-code Reader has a time offset.
5. A flag indicating whether the Time-code Reader has a frequency offset.
6. A flag indicating whether the timestamp came from the Time-code Reader. In case the Time-code Reader fails, the timestamp will be taken from the SBC system time. In this case the data are not used.
7. A flag indicating whether all ADC Boards had the same frame counter.
8. A flag indicating whether the Master ADC Board internal clock is working properly. This flag is not relevant in normal operation because an external clock is used instead.
9. A flag indicating whether the Master ADC Board is receiving an external 10-MHz clock.
10. A flag indicating whether the Master ADC Board is receiving an external 1-Hz clock.
11. Three bytes of padding, which contain no information, to make the total length a multiple of 4 bytes.

B.2 adc_sever bias Thread Actions

This section details the actions of the `adc_sever bias` thread, which depend on the request character received from `BiasServer`.

1. If the request character was ‘q,’ `adc_server` closed.
2. If the request character was ‘f,’ the thread sent the phase switching state information (4 bytes) immediately after reading it from VME address 0x10 on the Master ADC Board.
3. If the request character was ‘M,’ the thread sent the current bias bit values (2 bytes each for 16384 bias addresses¹) immediately after reading them from VME address 0xC00000 on the Master ADC Board.
4. If the request character was ‘m,’ the thread set the demodulation mask to blank the phase-switch transitions. The thread received the mask information (4 bytes) from `BiasServer`. The first two bytes specified the mask offset (typical value –1), and the second two specified the mask length (typical value 14). From this information, the thread constructed the 200-bit mask to send to the ADC Boards. Each bit indicated whether an 800-kHz sample was included in averaging to 4 kHz. A 1 indicated the corresponding sample was included, and a 0 indicated it was discarded.

The thread also constructed the quadrature mask, which indicated which sign to assign to each 800-kHz sample when constructing quadrature data. The quadrature mask was 90° out of phase with the demodulation. However, it must take into account the samples lost due to blanking. The offset of the quadrature mask (i.e. the location where the quadrature sign changed) was the offset of the demodulation mask plus

1. Not all addresses corresponded to a physical bias setting. There was no meaning to the values of the non-physical addresses.

half the demodulation mask length (because a longer demodulation mask shifted the center of each demodulation phase accounting for blanking) plus 1/4 of the total cycle length (to make it 90° out of phase)². The length of the quadrature mask (i.e. the length of each quadrature phase) was 100, half the total cycle. The thread wrote the demodulation mask to VME address 0x20000 on each ADC Board and the quadrature mask to address 0x60000 on each ADC Board.

5. If the request character was ‘z,’ then the thread sent the current demodulation mask offset and length (8 bytes total). These values were stored in `adc_server` memory and were not read from the ADC Boards. The `adc_server` set the mask on startup, and there was no other way to change the mask so the ADC Board value was always reflected in `adc_server` memory.
6. If the request character was ‘x,’ then the thread changed the set of housekeeping addresses to multiplex. First the thread received the number of addresses (4 bytes). Then the thread received the address values (2 bytes per address). Finally the thread wrote the address length to VME address 0xC on the Master ADC and the multiplex addresses to VME address 0x40000.
7. If the request character was ‘l,’ then the thread prepared a bias DAC command. It received the card, address, and DAC numbers; and the bias value (4 bytes total). It then wrote them to VME address 0x18 on the Master ADC. Because these commands often occurred in sequence, the thread paused after executing one in case the `fourhz_write` thread needed to communicate with the ADC Boards. The pause ensured that continued bias commands did not interrupt data recording.
8. If the request character was ‘b,’ then the thread instructed the Master ADC Board to send the bias DAC commands (prepared above) to the bias boards. The thread

2. A final offset of -1 was applied because the firmware applied a corresponding $+1$ offset.

received a mask (8 bytes total) indicating which cards, addresses, and DACs should be updated. The first bit was unused. The next 21 bits (corresponding to 21 card locations in the Bias-Board Backplane) indicated whether each card should be updated. Bits 24–63 indicated which addresses and DACs should be updated³. The thread wrote the mask information to VME addresses 0x1C (first 4 bytes) and 0x14 (second 4 bytes) on the Master ADC Board. The thread then waited for the Master ADC Board to complete the bias send. It indicated completion by setting the 8th bit (2^7 place) of the word at 0x0. The thread sent ‘q’ to BiasServer to indicate success.

9. If the request character was ‘p,’ then the thread commanded the Master ADC Board to set the phase-switch state. It received the setting (4 bytes) from the network and wrote it to VME address 0x10 on the Master ADC.
10. If the request character was ‘k,’ then the thread commanded the housekeeping multiplexing mode. It received 4 bytes from the network. If the first bit was 0, then the housekeeping data multiplexed addresses (see [6](#)). If the first bit was 1, then the housekeeping data recorded only a single address. The address was specified by the next 12 bits. In either case, the command was written to VME address 0x08 on the Master ADC.

B.3 data_serve Requests

This section lists the actions of the `DataCompilation` `data_serve` thread, which depended on the request received. Table [3.3](#) summarizes these requests.

1. If the request was `DATA_STATUS`, the thread sent the `DataCompilation` status flags (20 bytes, Appendix [A](#)).

3. These masks did not work properly. I recommend that the masks always be set to update all addresses and DACs. Failure to do so usually resulted in the bias boards receiving incorrect settings.

2. If the request was DATA_NADC, the thread sent the number of ADC Boards collecting data (§3.1.1).
3. If the request was DATA_DEMOD, the thread sent the latest radiometer data frame. The frame included the raw DE, TP, and quadrature data (after correcting the sign bit as described in §3.1.2), as well as the metadata: the number of 800-kHz samples accumulated into each 100-Hz sample, the 50-Hz and 1-Hz clock states, and firmware and `adc_server` status flags.
4. If the request was DATA_SNAP, the thread sent the latest snapshot.
5. If the request was DATA_HK, the thread sent the latest housekeeping data frame.
6. If the request was DATA_BIAS, the thread sent the biasing status. This included the bias board DAC bit values, the phase-switching state, and the demodulation mask.
7. If the request was DATA_PER, the thread sent the latest monitoring data from `PeripheralServer`. These included WPID settings, ADC Backplane monitoring data, CPID settings, cryostat pressure, and warm temperature sensing.
8. If the request was DATA_NEWSNAP, the thread signaled the `snap` thread with a semaphore to cause it to acquire a new snapshot. The `data_serve` thread then waited for the `snap` thread to respond by broadcasting a new snapshot condition (IEEE and The Open Group, 2012a). The thread then sent the request result (4 bytes⁴) and the identification number of the new snapshot (4 bytes).
9. If the request was DATA_ENCLOSURE, the thread sent the Electronics Enclosure smart power strip monitoring data. These included temperature, humidity, whether power output was enabled, and uptime.

4. Either an acknowledgment that the snapshot was updated or an error code indicating the reason for failure. See `DataInterfaceClient.h` in the RCS repository for the possible error codes.

10. If the request was DATA_PID_OUTPUT1 (2), the thread sent the status of the first (second) WPID output. The status was the fraction of the second during which the output was active i.e. driving the heater (fan).

B.4 DataCompilation–DataArc Communication

This section details the actions of `DataCompilation` when the `data_arc_service` thread received a message from `DataArc`. Table 3.4 summarizes the possible messages.

1. If the message was AGG_ACK, then `DataCompilation` received the frame number (4 bytes) from `DataArc`. If the frame number matched the current frame the thread sent, then the thread returned to waiting for a frame. By checking the frame number received by `DataArc`, `DataCompilation` never discarded data unless they were properly stored by `DataArc`.
2. If the message was AGG_FLUSH_BUFFER, then the thread removed all frames from the queue without sending them. This capability allowed `DataArc` to start a new data file with only new data⁵.
3. If the message was AGG_QUIT, `DataCompilation` closed the connection.
4. If the message was AGG_NACK, then `DataCompilation` received the frame number from `DataArc`. If the frame number was the same as the current frame, then the thread sent the frame.
5. In all other cases (including if there was no message after 50 ms), `DataCompilation` sent AGG_ACK_REQ followed by the current frame number.

5. The `DataCompilation` queue could contain arbitrarily old data.

B.5 ReceiverControl Commands

Table B.1 lists the commands `ReceiverControl` could process and the action it took for each.

Table B.1: `ReceiverControl` Commands

Command	Arguments	Description
PreampSet	<i>Diode Bias Offset</i>	Set the Preamp Board bias DAC corresponding to <i>Diode</i> to <i>Bias</i> bits. Then adjust the offset DAC with a feedback loop to center the channel output within the ADC range. <i>Offset</i> is an initial guess at the required offset DAC setting. Multiple commands can be issued in a single call by concatenating <i>Diode-Bias-Offset</i> triples with commas.
BiasAddr	<i>DAC Bits</i>	Use the Receiver Database to lookup the card, address, and DAC number of <i>DAC</i> . <i>Bits</i> may be optionally specified so that the result can be passed to bias-setting commands.
BiasIntAddr	<i>Card DAC Addr</i>	Convert the arguments to the corresponding bias address. The address is $DAC + 2^5 Addr + 2^9 Card$ and is the same as the address in Appendix B.2, bias command ‘M.’
SiteLookep	<i>Site</i>	Use the Receiver Database to find the module number currently located at electronics <i>Site</i> .
ModuleLookup	<i>Module</i>	Use the Receiver Database to find the electronics site where module number <i>Module</i> is located.
DbArc		Backup the content of the Receiver Database to disk.
BitSet	<i>DAC Bits</i>	Set bias electronics <i>DAC</i> output to <i>Bits</i> . Multiple commands can be issued by concatenating with commas.
AverageData	<i>Channel</i>	Print current average data in Volts.

Continued on Next Page...

Table B.1 – Continued

Command	Arguments	Description
MuxSet	<i>Addresses</i>	Change housekeeping multiplexing to monitor <i>Addresses</i> . Multiple addresses may be separated by commas, and addresses may contain MySQL wildcards. ReceiverControl will enlarge the address list with the ground and supply senses related to the input addresses, if applicable.
SetMask	<i>Offset Length</i>	Set the demodulation mask <i>Offset</i> and <i>Length</i> .
PhaseState	<i>States</i>	Change phase-switching state to <i>States</i> . Four states must be specified (the first two control one module leg; the second two control the other leg). Table 13 of Vanderlinde (2007) gives the meaning of the states.
AverageHk	<i>Frames</i>	Acquires Frames (250 ms each) of housekeeping data from DataCompilation . The number of frames is increased if necessary so that at least 1 full iteration of the housekeeping multiplexing is received. Print the resulting data, including conversion to physical units.
BitZeroNoPreamp	<i>DAC</i>	Set bias electronics to turn off modules except for preamp biasing (i.e. detector diodes will retain biasing). <i>DAC</i> may include MySQL wildcards. Multiple DACs may be concatenated with commas.
BitZeroYesPreamp	<i>DAC</i>	Same as BitZeroNoPreamp without restriction on preamp biasing.
Peripheral		See Dumoulin (2011) .
DrainVoltSet	<i>Drain Volts</i>	Adjust the module drain voltage of specified <i>Drain</i> to <i>Volts</i> . A feedback loop adjusts the corresponding gate voltages until the commanded drain voltage is reached. Multiple commands may be concatenated with commas. This command is only appropriate for Q band and W-band lab testing. W-band observation used MMIC bias boards with drain voltage control.

Continued on Next Page...

Table B.1 – Continued

Command	Arguments	Description
BiasArc	<i>Site</i>	Create a new entry in the Receiver Database with the current bias settings for modules matching <i>Site</i> , which may contain mySQL wildcards. Multiple <i>Sites</i> may be concatenated with commas.
ListBiasArc	<i>Module</i>	Print the entries in the Receiver Database containing stored bias settings for <i>Module</i> .
VerifyBiasArcUsed	<i>Id Module Tolerance</i>	Check that the current bias settings of <i>Module</i> match those from the Receiver Database record <i>Id</i> . <i>Tolerance</i> controls the allowed difference between the recorded and current housekeeping. Multiple tuples may be concatenated with commas.
UseBiasArcBit	<i>Id Module</i>	Set the bias settings of <i>Module</i> to those recorded in the Receiver Database record <i>Id</i> . Multiple tuples may be concatenated with commas.
UseBiasArcHk	<i>Id Module</i>	Set the bias settings of <i>Module</i> to those recorded in the Receiver Database record <i>Id</i> . The settings are adjusted so the current housekeeping matches the recorded values. Multiple tuples may be concatenated with commas.
StartSite	<i>Site Current1 Current2 Voltage1 Voltage2^a</i>	Turn on bias for modules at <i>Site</i> , which may contain mySQL wildcards. The first and second stage drain currents are set to <i>Current1</i> and <i>Current2</i> (specified in mA), respectively. The drain voltages are set to <i>Voltage1</i> and <i>Voltage2</i> in Volts. Multiple tuples may be concatenated with commas.
PhaseCurrSet	<i>Site A1Current A2Current B1Current B2Current</i>	Set phase switch currents of module at <i>Site</i> , which may contain mySQL wildcards. Currents are given in mA. Multiple tuples may be concatenated with commas.
PhaseBitSet	<i>Site A1Bits A2Bits B1Bits B2Bits</i>	Set DACs controlling phase switch currents of module at <i>Site</i> , which may contain mySQL wildcards. Multiple tuples may be concatenated with commas.

Continued on Next Page...

Table B.1 – Continued

Command	Arguments	Description
WhichSites	<i>Site</i>	Expand <i>site</i> , which may contain mySQL wildcards, into a list of unique sites. Multiple <i>Sites</i> may be concatenated with commas.
HktoDAC	<i>Address</i>	Print the identifier of the DAC corresponding to the given housekeeping <i>Address</i> . If there is no corresponding DAC, <i>Address</i> is printed. Multiple <i>Addresses</i> may be concatenated with commas.
MuxLookup	<i>NumericalAddress</i>	Print the address (name) of the housekeeping channel with the given <i>NumericalAddress</i> .
HKidLookup	<i>Address</i>	Print the numerical address of the housekeeping channel with the given <i>Address</i> identifier.
HKConvert	<i>Address Raw</i>	Convert housekeeping voltage <i>Raw</i> recorded for <i>Address</i> to physical units. Multiple tuples may be concatenated with commas.
HkConvertBatch	<i>Filename</i>	Convert raw housekeeping voltages from <i>Filename</i> to physical units.
DDLookup	<i>Diode</i>	Print ADC channel number recording data from detector <i>Diode</i> .
AdcChannelLookup	<i>Channel</i>	Print diode identifier corresponding to ADC <i>Channel</i> .
GetMask		Acquire the current demodulation mask from BiasServer and print it.
TakeSnapshot		Command DataCompilation to take a new snapshot and print the new snapshot number.
PerData	<i>Channel</i>	Print latest peripheral data from <i>Channel</i> . Multiple <i>Channels</i> may be concatenated with commas.
EnclosureData		Print latest monitoring data from the Electronics Enclosure smart power strip.
WPIDOutput		Print latest WPID output information.

^a Q-band modules require additional arguments *Current3* and *Voltage3*.

Appendix C

Timing Alignment Algorithm

`DataCompilation` aligns the radiometer and housekeeping frames:

1. Alignment of any sort is performed only for housekeeping and radiometer data. These are the only data which have GPS-based timestamps. For all other data (e.g. biasing), the most recent data frame available is used.
2. Frame construction waits until the queue of radiometer data and the queue of housekeeping data each have at least 4 250-ms frames.
3. Compute `TIMING_FLAG` to be true if the Time-code Reader is synchronized to the IRIG signal and the firmware is synchronized to the time-code reader and the timestamp source was the Time-code Reader. Otherwise, `TIMING_FLAG` is false.
4. If the timestamp of the front frame in the radiometer queue is not $0 \pmod{250 \text{ ms}}$ or `TIMING_FLAG` is false, the frame is moved into the first radiometer 250-ms frame of the 1-s frame.
5. Otherwise if the timestamp of the front frame is $0 \pmod{250 \text{ ms}}$ but not $250 \text{ ms} \pmod{1 \text{ s}}$, a null 250-ms frame is copied into the 1-s frame. The front frame remains in the queue.
6. The previous two steps are repeated with “radiometer” replaced by “housekeeping.”
7. If the first radiometer 250-ms frame is empty, and the first housekeeping 250-ms frame is non-empty (or vice versa), the front radiometer (or housekeeping) 250-ms frame is moved to the 1-s frame.
8. If both first 250-ms frames in the 1-s frame are empty, and the two front 250-ms frames have the same timestamp, both front 250-ms frames are moved into the 1-s frame.

9. If both first 250-ms frames are empty, and the two front 250-ms frames do not have the same timestamp, the 250-ms frame with the earlier timestamp is moved, and a null 250-ms frame is copied into the other 250-ms frame. The front 250-ms frame with later timestamp remains in the queue.
10. The preceding 5 steps are repeated with “first 250-ms frame” replaced by “second,” “third,” and “fourth.” “not 250 ms” is replaced by 500 ms, 750 ms, and 0 ms.
11. The latest non-null timestamp¹ is copied to frame.utc (the value cbicontrol uses to align receiver and mount data).
12. If TIMING_FLAG is true and the first housekeeping 250-ms frame’s timestamp is 250 ms (mod 1 s) and each housekeeping 250-ms frame has the same timestamp as the corresponding radiometer 250-ms frame and the radiometer 250-ms frames’ timestamps increment by 250 ms and frame.utc is 1 s greater than the previous value of frame.utc, then increment the good timing count. Otherwise, set the good timing count to 0.
13. If the good timing count is at least 10, request (by setting the time alignment register) that cbicontrol perform strict alignment using this frame’s timestamp. Otherwise, request that cbicontrol ignore this frame’s timestamp and put this frame into the earliest available multi-system frame.
14. Move this frame into the buffer of frames to send to cbicontrol, and return to step 2 to construct the next frame.

1. This algorithm guarantees at least 1 non-null 250-ms frame.

Appendix D

Receiver Database

This appendix summarizes the information in the Receiver Database and the mapping between modules, MABs, and bias boards. The Receiver Database was a mySQL database that stored the relationships between the physical configuration of the instrument and the electronics numbering. The mySQL table ADC_MAPPING (Appendix D.1) stored the relationship between modules and ADCs. The mySQL table BIASES (Appendix D.2) stored the relationship between a bias voltage or current and the DAC that provided it. The mySQL table HOUSEKEEPING (Appendix D.3) stored the relationship between housekeeping addresses and physical monitored items. The mySQL table mab_input (Appendix D.4) stored which modules were in which MABs. The lists in the following sections give the meanings of the columns in the mySQL tables.

D.1 Table ADC_MAPPING

Each row in table ADC_MAPPING corresponded to 1 diode.

- NAME is the diode name arranged by MAB e.g. M08A6Q2 is the Q2 diode of the module at site 6 of MAB 8.
- MODULE is the module serial number.
- TYPE is W1 or W2. W1 means that the module has two gate controls for the first MMIC. W2 means that the module has 1 gate control per MMIC.
- ADC_CH is the channel number (ADC) recording the diode.
- PR_SLOT is the location of the Preamp Board (in the Bias-Board Backplane) for that diode.

- PR_SIDE is the side (0 or 1) of the Preamp Board used.
- ADC_SLOT is the number (1–13) of the ADC Board recording that diode.
- PR_SERIAL is the Preamp Board serial number.
- ADC_SERIAL is the ADC Board serial number.

D.2 Table BIASES

Each row in table BIASES corresponded to 1 bias DAC.

- DESCRIPTION is the identifier (e.g. M12A6Q1B) of a bias setting.
- SLOT is the location in the Bias-Board Backplane of the Bias Board controlling that setting.
- DAC is the number of the DAC chip on that Bias Board controlling the bias setting.
- ADDR is the “address” (1–8) of the DAC on the DAC chip.
- INT_ADDR is firmware bias memory location storing the bias setting.

D.3 Table HOUSEKEEPING

Each row in table HOUSEKEEPING corresponded to 1 monitored item.

- NAME is the identifier (e.g. M12A6DA2I) of the voltage, current, temperature, etc. being monitored.
- DESCRIPTION identifies the locations of thermometers.
- MUX_INT is the multiplex address for the identifier.
- MUX is a string holding the binary digits of MUX_INT.

- TYPE identifies the type of measurement:
 1. GN is a ground voltage.
 2. G is a gate voltage.
 3. DI1 is a first-stage drain current.
 4. DV1 is a first-stage drain voltage.
 5. DI2 is a second-stage drain current.
 6. DV2 is a second-stage drain voltage.
 7. S25 is a 2.5-V supply voltage.
 8. P is a phase-switch current.
 9. CD is a cryogenic thermometer.
 10. CD4 is a LakeShore DT-470 cryogenic thermometer.
 11. CD6 is a LakeShore DT-670 cryogenic thermometer.
 12. TH is a warm (Electronics Enclosure) thermometer.
- BOARD_SLOT is the location in the Bias-Board Backplane of the Bias Board controlling the voltage or current being monitored.
- BOARD_SIDE is the side of the Bias Board.
- BOARD_SN is the serial number of the Bias Board.

D.4 Table mab_input

Table D.1 lists the bias boards connected to each MAB. Table D.2 lists the modules on each MAB and the ADC Board digitizing each.

Table D.1. MAB–bias-board Mapping

MAB	MMIC Board Location/Side	Phase-switch Board	Preamp Board
0	14/1	20/1	5/1
1	15/0	20/2	6/0
2	14/0	20/0	5/0
3	15/1	21/0	6/1
4	12/0	18/1	3/0
5	16/0	21/1	7/0
6	11/1	18/0	2/1
7	13/0	19/1	4/0
8	16/1	21/2	7/1
9	13/1	19/2	4/1
10	10/0	17/0	1/0
11	11/0	17/2	2/0
12	10/1	17/1	1/1

Table D.2. Modules and ADC Board for Each MAB

MAB	Modules (RQ)	ADC Board
0	0, 1, 2, 3, 4, 5, 9	9
1	10, 11, 12, 17, 18, 19, 20	10
2	6, 7, 8, 13, 14, 15, 16	8
3	24, 25, 26, 27, 28, 29, 36	11
4	21, 22, 23, 30, 31, 32, 33	5
5	37, 38, 39, 47, 48, 49, 50	12
6	34, 35, 44, 45, 46, 55, 45	4
7	40, 41, 42, 43, 51, 52, 53	6
8	57, 58, 59, 60, 67, 68, 69	13
9	54, 61, 62, 63, 64, 65, 66	7
10	74, 75, 76, 77, 82, 83, 84	1
11	70, 71, 72, 73, 78, 79, 80	3
12	85, 86, 87, 88, 89, 90	2

Appendix E

Database and Data-quality Statistics

This appendix summarizes the QUIET mySQL Database, including the data-quality statistics for each CES-diode. Database `quiet_cal` (Appendix [E.1](#)) contained calibration information, particularly the responsivity model. Database `quiet_data` (Appendix [E.2](#)) contained information about observations and data files. Database `quiet_quality` (Appendix [E.3](#)) contained the data-quality statistics we calculated for each CES-diode. The following sections list the meanings of each column in each mySQL table.

E.1 Database `quiet_cal`

This database contained calibration information.

E.1.1 Table `w91_gain_model_each_runs`

Table `w91_gain_model_each_runs` contained the latest responsivity model, currently identical to `w91_gain_model_each_run_2012_0404`. Each row contained the responsivity for 1 CES-diode.

- `ces_id` is a unique identification number for each CES, used internally in the database.
- `run_id` is the CES run number.
- `run_subid` is the number of the CES within its run. Together `run_id` and `run_subid` uniquely specify the CES.
- `module` is the module (RQ) number.
- `diode` is Q1, U1, U2, or Q2.

- version was used to distinguish different models in the same table; however, the final table contains only 1 model (version 0).
- G is the CES-diode responsivity in mV/K, thermodynamic temperature.

E.2 Database quiet_data

This database contained observation, data-file, and log information. Table w91_ces contained the definition of each CES. Table w91_l0_ces_intersection listed the data files for each CES. Table w91_level0 contained summary information about each Level-0 data file. Table w91_level1 contained summary information about each Level-1 data file. Table w91_run_log contained details about each observation run.

E.2.1 Table w91_ces

Each row contained the definition of 1 CES.

- ces_id, run_id, run_subid are as above.
- begin_time is the start of the CES in UTC. We included this second.
- end_time is the end of the CES in UTC. We did not include this second.

E.2.2 Table w91_l0_ces_intersection

This table contained the relationship between CESes and Level-0 data files. Each row corresponded to the fraction of 1 Level-0 file containing data for 1 CES.

- l0_id The identification number of the Level-0 file.
- ces_id is as above.
- begin_time is the number of minutes into the Level-0 file when the CES begins.

- `end_time` is the number of minutes into the Level-0 when the CES ends.
- `l0_length` is the length of the Level-0 file in minutes.

E.2.3 Table `w91_level0`

Each row contained information about 1 Level-0 file.

- `id` is the Level-0 file identification number.
- `path` is the filename including directories from the data storage root location.
- `type` is one of
 1. RCS indicating the old Level-0 format recorded by `DataArc`.
 2. ARCHIVE indicating normal Level-0 files created by `cbicontrol`.
 3. ARC_PRE indicating old Level-0 files created by `cbicontrol` but not including receiver data.
- `hash` is the SHA-1 checksum of the file.
- `transfer` indicates the data transfer status:
 1. DVD means we recorded the file to Blu-ray disk.
 2. SHIP means we sent the disk to Chicago.
 3. FAIL means the file was corrupt on Blu-ray disk and we need to send it again.
 4. COMPL means we transferred the file successfully.
- `available` indicates where the file is available
 1. SITE means it is available at the site storage.
 2. CHI means it is available at the Chicago data center.

3. KEK means it is available at the KEK data center.
4. OSLO means it is available at the Oslo data center.
5. CHI_TAPE means it is available on tape archive at Chicago.

- begin_time is the start time of the file in UTC.
- end_time is the end time of the file in UTC.
- begin_db_ver is the Receiver Database version in use at the beginning of the file.
- begin_db_hash is the checksum of the Receiver Database in use at the beginning of the file.
- end_db_ver is the Receiver Database version in use and the end of the file.
- end_db_hash is the checksum of the Receiver Database in use at the end of the file.
- comment contains comments about each file.

E.2.4 Table w91_level1

Each row contained information about 1 Level-1 file.

- id is the identification number of the Level-1 file.
- path is as above.
- hash is as above.
- l0_id is the identification number of the corresponding Level-0 file.
- available is as above
- flag indicates several problems:

1. CORRUPT_FILE means the file cannot be read.
2. CORRUPT_SUBTIME means the receiver timestamp in the file is impossible e.g. before the experiment started.
3. WRONG_SUBTIME means the receiver timestamp does not match the expected time when the file was recorded.

- num_hdu is the number of FITS HDUs in the file. Valid files have 1876 HDUs.

E.2.5 Table w91_run.log

Each row contained information about 1 observation run.

- id is the identification number of the run.
- begin_time is the start of the run in UTC.
- end_time is the end of the run in UTC.
- run_type is:
 1. CMB for CMB patch observations (CESes).
 2. Calib for calibration observations.
 3. Other for all other observations.
- object is the astronomical source.
- state is
 1. Done means the observation finished successfully.
 2. Scheduled means we scheduled the observation but have not confirmed it.
 3. Canceled means we canceled the observation before it started.

4. Aborted means we stopped the observation before it completed.

- `daily_log_scheduled` is the identification number of the observer log when when scheduled the observation.
- `daily_log_confirmed` is the identification number of the observer log when we confirmed the observation completed.
- `comment` contains comments for each run. Usually we put additional details such as the deck angle, modules, or details of special calibrations.

E.3 Database `quiet_quality`

This database contained the data-quality statistics for each CES-diode. Table `w91_atm` contained weather information. Table `w91_ces_usable` defined whether each CES was usable based on telescope, mount, and software problems. Table `w91_diode_usable` defined whether each CES-diode was usable. Table `w91_diodes` contained information about each diode. Table `w91_full_correlation` contained correlation matrices for the noise model. Table `w91_housekeeping` contained average housekeeping for each CES. Table `w91_scan` contained summarized pointing information for each CES. Table `w91_software` contained details about software problems for each CES. Table `w91_sun_spike` defined the Sun-sidelobe cut. Table `w91_timestream` contained data-quality statistics for each CES-diode. Table `w91_typeb` contained Type-B-glitching statistics for each CES-diode.

E.3.1 Table `w91_atm`

Each row contained weather information for 1 CES.

- `ces_id`, `run_id`, and `run_subid` are as above
- `humidity_med` is the median humidity during the CES in %.

- `humidity_min` is the minimum humidity during the CES in %.
- `humidity_max` is the maximum humidity during the CES in %.
- `humidity_rms` is the RMS fluctuation of the humidity during the CES in %.
- `air_temp_med`, `air_temp_min`, `air_temp_max`, `air_temp_rms` are the same for the air temperature in °C.
- `pressure_*` are the same for the atmospheric pressure in hPa.
- `wind_speed_*` are the same for the wind speed in m/s.
- `wind_dir_*` are the same for the wind direction in °.
- `apex_humidity_*` are the same for the APEX weather station.
- `apex_dewpoint_*` are the same for the dew point in °C from APEX.
- `apex_pwv_*` are the same for the PWV in mm from APEX.

E.3.2 Table `w91_ces_usable`

Each row defines whether 1 CES was usable.

- `run_id` and `run_subid` are as above.
- `usable` indicates whether the CES is usable.
- `usable_daq` indicates whether the CES is usable based on software problems.
- `usable_tel` indicates whether the CES is usable based on telescope and mount problems.

E.3.3 Table w91_diode_usable

Each row defined whether 1 CES-diode was usable. To be usable, both w91_ces_usable.usable and w91_diode_usable.usable must be true.

- ces_id, run_id, run_subid, module, diode are as above.
- usable indicates whether the CES-diode is usable.

E.3.4 Table w91_diodes

Each row contained information on 1 diode.

- module and diode are as above
- knee_median_Hz is the median knee frequency in Hz of the diode for all CESes passing data selection.
- knee_mean_Hz is the mean knee frequency.
- knee_RMS_Hz is the RMS fluctuation of the knee frequency.
- uK_sqrtsec_* are the same for the sensitivity in $\mu\text{K}\sqrt{\text{s}}$
- CES_count is the number of CESes passing data selection.
- leakage is the $I \rightarrow Q/U$ leakage in %.
- MAB is the MAB that diode is on.
- gain_mean is the mean responsivity for all CESes passing data selection.
- band_center is the bandpass center frequency in GHz.
- band_pass is the bandwidth in GHz.

E.3.5 Table w91_full_correlation

Each row contained 1 element of the white and 1/f-noise correlation matrices.

- run_id and run_subid are as above.
- module1 is one of the modules whose noise correlation is measured.
- diode1 is one of the diodes whose noise correlation is measured.
- module2 and diode2 specify the other diode.
- correl_oneoverf is the 1/f-noise correlation coefficient.
- correl_white_noise is the white-noise correlation coefficient.

E.3.6 Table w91_housekeeping

Each row contained housekeeping information for 1 CES.

- run_id and run_subid are as above.
- preamp1_temp_p3t4_mean is the mean temperature in °C of a thermometer on a Preamp Board.
- preamp1_temp_p3t4_rms is the RMS temperature fluctuation during the CES.
- preamp1_temp_p3t4_range is the range (maximum – minimum) during the CES.
- mmic3_temp_p3t9_* are the same for a MMIC Board.
- psw1_temp_p3t10_* are the same for a Phase-Switch Board.
- psw2_temp_p3t11_* are the same for a second Phase-Switch Board.

- `primary_mirror_temp_0_*` are the temperatures in K for a thermometer on the primary mirror.
- `secondary_mirror_temp_3_*` are the same for the secondary mirror.
- `groundscreen_temp_7_*` are the same for the ground screen.
- `external_air_temp_26_*` are the same for a thermometer exposed to the ambient environment.
- `module_45_temp_p2t19_*` are the same for a thermometer on module RQ45.
- `80K_temp_p2t18_*` are the same for the 80-K radiation shield.
- `platelet1_temp_p2t16_*` are the same for the platelet array.
- `interface1_temp_p2t14_*` are the same for the interface plate.
- `module_9_temp_p2t8_*` are the same for RQ9.

E.3.7 Table w91_scan

Each row contained pointing information for 1 CES.

- `run_id` and `run_subid` are as above.
- `azimuth` is the average azimuth of the CES in $^{\circ}$ East of North.
- `elevation` is the average elevation of the CES.
- `deck_angle` is the deck angle of the CES.
- `scan_speed` is the average azimuthal scan speed in $^{\circ}/s$.
- `scan_period` is the scan period in s.

- time is the time in days since an arbitrary zero point.
- object is the patch name.
- duration is the CES length in s.
- sun_ra is the RA of the Sun at the time of the CES.
- sun_dec is the Declination of the Sun at the time of the CES.
- sun_phi is the azimuthal angle of the Sun in the telescope-centered coordinate system.
- sun_theta is the distance to the Sun in $^{\circ}$.
- moon_* are the same for the Moon.
- hour_angle is the hour angle in $^{\circ}$.
- patch_day_number is a number for dividing CESes into days for day-by-day cross-correlation. If two CESes scan the same patch, then they have the same patch_day_number iff the patch did not set between the two CESes.
- hour_of_day is the hour into the (UTC) day when the observation was performed.
- sidelobe_az is the azimuth at which the triple-reflection sidelobe is pointing.
- sidelobe_el is the elevation at which the triple-reflection sidelobe is pointing.
- sidelobe2_* are the same for the spillover sidelobe.
- array_orientation_equatorial is the orientation axis of the array on the sky in Equatorial coordinates.
- scan_direction_equatorial is the scan axis on the sky in Equatorial coordinates.
- sidelobe_lat is the latitude of the triple-reflection sidelobe in Galactic coordinates.

- `sidelobe_long` is the longitude of the triple-reflection sidelobe in Galactic coordinates.
- `sun_az` is the azimuth of the Sun at the time of CES observation.
- `sun_el` is the elevation of the Sun at the time of CES observation.

E.3.8 Table w91_software

Each row contained software status for 1 CES.

- `run_id` and `run_subid` are as above.
- `processing_status` indicates possible failures when the CES was being checked. The possible values are
 1. `GOOD` indicates no problems occurred
 2. `NO_TELESCOPE` indicates the `TELESCOPE` HDU was missing or could not be read from at least 1 of the Level-1 files constituting the CES.
 3. `PARTIAL_FRAME` indicates that the number of telescope rows read was not a multiple of 100 (i.e. not an integral length in seconds).
 4. `EXTRA_FILE` indicates a Level-1 file was in the CES intersection but no corresponding rows could be found in it.
 5. `CES_LOOKUP_FAIL` indicates the CES metadata could not be found. This is more likely to be a problem with the checking code than with the CES itself.
- `num_frame` is the number of 1-Hz frames in the CES (i.e. the length in seconds).
- `num_uneven_demod` is the number of uneven demodulation events detected in the ADC. This must be 0 if the CES can be used.
- `num_telescope_time_jump` is the number of telescope timestamps with irregular intervals.

- `num_uneven_mask` is the number of mask mismatch events between ADC boards. This is contaminated by snapshot bugs.
- `num_uneven_clock` is the number of fast clock mismatch events between ADC boards. This is contaminated by snapshot bugs.
- `num_bad_scale` is the number of bad scale factor events. This must be 0.
- `num_bad_quad_scale` is the number of bad quadrature scale factor events. This must be 0.
- `num_4hz_mismatch` is the number of 4-Hz frame counter mismatch events across ADC boards. This must be 0 unless the ADC was started incorrectly.
- `num_tfp_time_offset` is the number of “time offset” events detected from the Time-code Reader.
- `num_tfp_freq_offset` is the number of “frequency offset” events detected from the Time-code Reader.
- `num_null_frame` is the number of null frames in the data. This must be 0.
- `num_bad_timing_flag` is the number of frames flagged as having bad timing by the Online Software. This must be 0.
- `num_bad_wpid_thread` is the number of frames when the WPID output was not being monitored.
- `num_rcs_frame_jump` is the number of frame gaps due to missing data lost between the Online Software and data being written to disk. This must be 0.
- `num_rcs_frame_offset` is the number of frames whose timestamps are not aligned to the start of the second. This must be 0.

- `num_adc_timestamp_differ` is the number of frames having different timestamps between radiometer and housekeeping. This must be 0.
- `num_adc_timestamp_jump` is the number of irregular radiometer timestamps. This must be 0.
- `num_hk_frame_jump` is the number of frame gaps due to missing housekeeping data. This must be 0.
- `num_adc_frame_jump` is the number of frame gaps due to missing radiometer data. This must be 0.
- `fiftyhz_clock_bad` indicates whether a problem was found with the ADC Board 50-Hz clock. This must be 0.
- `onehz_clock_bad` indicates whether a problem was found with the ADC Board 1-Hz clock. This must be 0.
- `psw_state*` are the phase-switch switching state parameters.
- `num_bad_mask` is the number of bad blanking mask events. This must be 0.
- `num_adc_count_mismatch` is the number of frames with different radiometer and housekeeping counters. This must be 0.
- `num_bad_mux_pattern` is the number of housekeeping samples with irregular multiplexing. This must be 0.
- `num_mux_addr` is the number of different housekeeping addresses sampled. This must be 1774.
- `num_downsample_frame` is the number of frames when the Online Software was operating in $10 \rightarrow 1$ downsampling mode. This must be 0.

- `num_bad_psw_flag` is the number of events indicating the phase switching state was not known.
- `num_telescope_radiometer_mismatch` is the number of samples where the telescope (mount) and radiometer timestamps do not match. This must be 0.
- `wrong_len` indicates whether the different Level-1 HDUs had compatible lengths. This must be 0.
- `max_telescope_time_jump` is the maximum irregularity in the telescope timestamps (in ms). This must be < 5 .

E.3.9 Table `w91_sun_spike`

Each row defined whether 1 CES-module was rejected by the Sun-sidelobe cut.

- `run_id`, `run_subid`, and `module` are as above.
- `has_spike` is non-zero if the CES-module should be rejected because of Sun-sidelobe contamination.

E.3.10 Table `w91_timestream`

Each row contained data-quality statistics for 1 CES-diode.

- `run_id`, `run_subid`, `module`, and `diode` are as above.
- `uK_sqrtsec` is the white-noise level in $\mu\text{K}\sqrt{\text{s}}$.
- `knee_frequency` is the knee frequency in Hz.
- `spectral_index` is the 1/f slope (in noise power).
- `fit_flag` indicates problems in the noise fit:

- 0 means fit success
- 1 means the fitter reached the limit for the number of likelihood function evaluations
- 10 means the 1/f knee frequency was forced to 0 and the slope to 1.2
- `scan_synchronous_signal2` is the ratio of the noise power at the scan frequency to the power predicted by the noise model.
- `scan_synchronous_signal` is the statistical significance (σ) of `scan_synchronous_signal2`.
- `scan_synchronous_signal2_averg` is the same as `scan_synchronous_signal2`, but computed from TP instead of DD.
- `scan_synchronous_chi2` is the reduced χ^2 to 0 of the signal removed by the ground subtraction filter.
- `excess_10mHz_100mHz` is the average difference between the noise power and model in σ s between 10–100 mHz.
- `excess_100mHz_2Hz` is the same between 100 mHz and 2 Hz.
- `chisquare_10mHz_100mHz` is the reduced χ^2 between the data and noise model between 10–100 mHz.
- `chisquare_excl_scan_10mHz_200mHz` is the same as `chisquare_10mHz_100mHz` except we excluded the scan frequency.
- `chisquare_100mHz_2Hz` is the same as `chisquare_10mHz_100mHz` except between 100 mHz and 2 Hz.
- `binned_spectrum_spike_sigma` is the maximum difference between the data noise spectrum and the model in σ .

- `binned_spectrum_spike_frequency` is the frequency of the maximum. I limited the search to the range between twice the scan frequency and 9.6 Hz which the filter accepts.
- `sss_filtered_binned_spectrum_spike_sigma` is the same as `binned_spectrum_spike_sigma` except I applied the ground subtraction filter first.
- `sss_filtered_binned_spectrum_spike_sigma_1.2Hz` is the same as `sss_filtered_binned_spectrum_spike_sigma` except the search range is 1.15–1.3 Hz.
- `glitch_1sample` is the largest 1-sample (20-ms) deviation from the mean in the CES-diode in σ .
- `glitch_100msec` is the largest deviation of 100-ms moving averages.
- `glitch2_1sample` is the same as `glitch_1sample` except we used the data after filtering.
- `jump_demod_1sec` is the largest change in the 1-s averaged DD data.
- `jump_averg_1sec` is the same as `jump_demod_1sec`, except for TP data.
- `std_of_std_10sec` is the fractional variation in the noise level during the CES computed from the standard deviation of the standard deviations of 10-s data chunks.
- `fiducial_gain` is the responsivity in mV/K used in processing for this table. It can be different from the final responsivity for CMB analysis.
- `filtered_chisquare_10mHz_200mHz` is the same as `chisquare_10mHz_200mHz` except we applied a frequency-domain bandpass filter (not the same as the fiducial analysis filter).
- `sss_filtered_chisquare_nearscan` is the reduced χ^2 for the noise model fit for the 40 Fourier modes nearest the scan frequency. We applied the ground subtraction filter first.

- `each_az_chisquare_below_1Hz` is the reduced χ^2 for the noise model fit below 1 Hz. We applied the full analysis filter first.
- `each_az_dof_below_1Hz` is the number of degrees of freedom for `each_az_chisquare_below_1Hz`.
- `each_az_chisquare_slopes` is the χ^2 to 0 for the azimuth slopes subtracted in filtering.
- `five_min_max_raw` is the maximum subtracted slope ($\mu\text{K}/^\circ$) in a 150-s moving average.
- `five_min_max_normalized` is `five_min_max_raw` times the width of the scan in azimuth i.e. the maximum signal change in μK attributed to the slope.

E.3.11 Table w91_typeb

Each row contained Type-B statistics for 1 CES-diode.

- `run_id`, `run_subid`, `module`, and `diode` are as above.
- `chisquare_demodav_uncorrected` is the Type-B χ^2 before Type-B-glitch correction.
- `chisquare_demodav_corrected` is the Type-B χ^2 after Type-B-glitch correction.
- `bits_sqrtsec_uncorrected` is the noise level in ADC bit $\sqrt{\text{s}}$ units before Type-B correction.
- `bits_sqrtsec_corrected` is the noise level found in iterative correction.
- `rms_1sec_uncorrected_volts` is the RMS fluctuation in V after creating 1-s samples by averaging.
- `min_glitch_distance_volts` is the minimum distance to any of the Type-B glitches.
- `mean_tp_plus_volts` is the average TP for one of the 50-Hz phase-switch states.
- `mean_tp_minus_volts` is the average TP for the other phase-switch state.

- `mean_demod_plus_corrected_volts` is the mean DE for one of the 50-Hz phase-switch states.

Appendix F

Array Sensitivity Definition and Calculation

This appendix defines the “QUIET array sensitivity” in a way that takes the white noise correlations within each module into account. The definition has several steps. See [Buder \(2010c\)](#); [Kusaka et al. \(2012\)](#) for more detail.

F.1 Diode Sensitivity

I defined the sensitivity of each diode as its mean white noise level, averaged over those CESes for which the diode passed data selection.

F.2 Diode-pair Sensitivity

A “diode pair” is two Q diodes or two U diodes in the same module. The white noise in a diode pair tends to have positive correlation, so the sensitivity to Q or U of a diode pair is better than one would expect if one ignored the correlation and treated the diodes as uncorrelated. Because the detector angles in a diode pair are nearly 90° apart, we effectively differenced the diodes in map making. We weighted diodes by their inverse variance. Therefore, as an approximation to the sensitivity in the map, I defined the diode-pair sensitivity to be the standard deviation of the inverse-variance-weighted diode difference,

$$d = (Q_1 w_1 - Q_2 w_2) / w, \quad (\text{F.1})$$

where Q_i is the data of one of the diodes,

$$w_i = 1/\sigma_i^2 \quad (\text{F.2})$$

are inverse variance weights,

$$w = w_1 + w_2 \quad (\text{F.3})$$

is the weight total, and σ_i is the sensitivity of diode i . The variance is

$$\langle d^2 \rangle = \frac{\langle Q_1^2 \rangle w_1^2 + \langle Q_2^2 \rangle w_2^2 - 2w_1 w_2 \langle Q_1 Q_2 \rangle}{w^2}. \quad (\text{F.4})$$

Neglecting any signal, $\langle Q_i^2 \rangle = \sigma_i^2$. Therefore

$$\langle d^2 \rangle = \frac{w_1 + w_2 - 2\rho_{12}\sigma_1^{-1}\sigma_2^{-1}}{w^2}, \quad (\text{F.5})$$

where ρ_{12} is the white noise correlation coefficient. I used the mean correlation coefficient from CESes where the Q1 or U1 diode of a pair passed data selection. (The difference if I used CESes where Q2 or U2 passed was $\lesssim 0.001$.) The diode-pair sensitivity is $\sigma_{12} = \sqrt{\langle d^2 \rangle}$. In case one of the diodes in a pair is non-functional, the diode pair sensitivity is the sensitivity of the functional diode.

F.3 Module Sensitivity

Since the diode pairs in a module were nearly uncorrelated (Bischoff & Kusaka, 2009), I treated the module's Q and U sensitivities as independent. I defined the combined sensitivity to polarization in a module as

$$\sigma = 1/\sqrt{\sigma_Q^{-2} + \sigma_U^{-2}}, \quad (\text{F.6})$$

where σ_Q is the sensitivity of the Q diode pair in the module.

F.4 Array Sensitivity

The correlations between modules were negligible, so I combined modules in the same way I combined diode pairs in a module. The array sensitivity (including all modules that passed data selection at least once) was $87 \pm 7 \mu\text{K}\sqrt{\text{s}}$ ([Kusaka et al., 2012](#)). The mean module sensitivity was $899 \mu\text{K}\sqrt{\text{s}}$. I defined an “effective average module sensitivity,” $756 \mu\text{K}\sqrt{\text{s}}$, to be the array sensitivity times the square root of the number of modules. An array of modules each having this sensitivity would have the same sensitivity as the real array.

Appendix G

Null-test Definitions and Correlations

Here I define each of the null-test divisions. See also [Buder et al. \(2012\)](#). See §5.6.2 for general discussion of the null tests.

Table G.1: Null-test Definitions

Name	Description	Database Query
typeo_pointing ^a	CES in Type-O Effect pointing region vs. others	(((w91_scan.deck_angle BETWEEN -200. AND -100.) OR (w91_scan.deck_angle BETWEEN 160. AND 260.)) AND (w91_scan.elevation BETWEEN 65. AND 75.))
I_to_Q_by_diode	Diodes with large vs. small $I \rightarrow Q$ leakage	(ABS(w91_diodes.leakage) > 0.479999989271164)
Q_vs_U_diode	Q diodes vs. U diodes	(w91_timestream.diode = 'Q1') OR (w91_timestream.diode = 'Q2')
Q1U1_vs_Q2U2_diode ^b	Q1 and U1 diodes vs. Q2 and U2 diodes	(w91_timestream.diode = 'Q1') OR (w91_timestream.diode = 'U1')
Q1U2_vs_Q2U1_diode ^b	Q1 and U2 diodes vs. Q2 and U1 diodes	(w91_timestream.diode = 'Q1') OR (w91_timestream.diode = 'U2')
band_cf_low_vs_high	Module band-pass center frequency low vs. high	(w91_diodes.band_center <= 93.8)
array_top_vs_bottom	Modules at the array top vs. bottom	(w91_timestream.module <= 39)
array_central_vs_peripheral	Central vs. peripheral modules	(w91_timestream.module BETWEEN 15 AND 18 OR w91_timestream.module BETWEEN 23 AND 27 OR w91_timestream.module BETWEEN 32 AND 37 OR w91_timestream.module BETWEEN 42 AND 48 OR w91_timestream.module BETWEEN 53 AND 58 OR w91_timestream.module BETWEEN 63 AND 67 OR w91_timestream.module BETWEEN 72 AND 75)

Continued on Next Page...

Table G.1 – Continued

Name	Description	Database Query
typeb_happening_vs_not_happening	Type-B glitching happening vs. not happening	(ABS(w91_typeb.chisquare_demodav_uncorrected - w91_typeb.chisquare_demodav_corrected) > 0.1)
sss_large_vs_small	Scan synchronous signal large vs. small	(w91_timestream.scan_synchronous_signal > 0.602676) ^c
high_freq_excess_large_vs_small	High frequency noise spikes vs. no spikes	(w91_timestream.chisquare_2Hz_15Hz > 1.244210) ^c
elevation_high_vs_low	Elevation high vs. low	(w91_scan.elevation > 65.150700) ^c
sun_close_vs_far	Sun close vs. far to the main beam	(w91_scan.sun_theta < 69.490200) ^c
moon_close_vs_far	Moon close vs. far	(w91_scan.moon_theta < 97.029000) ^c
enclosure_temp_high_vs_low	Electronics enclosure temperature high vs. low	(w91_housekeeping.preamp1_temp_p3t4_mean > 45.595500) ^c
enclosure_temp_change_large_vs_small	Enclosure temperature change within the CES large vs. small	(w91_housekeeping.preamp1_temp_p3t4_range > 0.707993) ^c
cryostat_temp_change_large_vs_small	Cryostat temperature change large vs. small	(w91_housekeeping.module_45_temp_p2t19_range > 0.707127) ^c
weather_bad_vs_good	Bad weather vs. good	(w91_weather.tp_rms_10sec > 0.098875) ^c
pwv_high_vs_low	PWV high vs. low	(w91_atm.apex_pwv_med < 0.800000) ^c
1sthalf_vs_2ndhalf	First half of the season vs. second	(w91_scan.time <= 611)
knee_large_vs_small	Diodes with 1/f knee frequency large vs. small	(SELECT AVG(w91_diodes_knee_temp.knee_median_Hz) FROM w91_diodes AS w91_diodes_knee_temp WHERE w91_diodes_knee_temp.knee_median_Hz > 0 AND w91_diodes_knee_temp.module = w91_timestream.module) > 0.00862041371874511)
sensitivity_high_vs_low	Sensitivity high vs. low diodes	(w91_diodes.uK_sqrtsec_median > 1494.25)

Continued on Next Page...

Table G.1 – Continued

Name	Description			Database Query	
deck_phase0_90	Deck angles apart	90°		(ROUND(w91_scan.deck_angle)=30 ROUND(w91_scan.deck_angle)=75 ROUND(w91_scan.deck_angle)=210 ROUND(w91_scan.deck_angle)=-105 ROUND(w91_scan.deck_angle)=-150)	OR OR OR OR OR
deck_phase0_45	Deck angles apart	45°		(ROUND(w91_scan.deck_angle)=30 ROUND(w91_scan.deck_angle)=120 ROUND(w91_scan.deck_angle)=210 ROUND(w91_scan.deck_angle)=-60 ROUND(w91_scan.deck_angle)=-150 ROUND(w91_scan.deck_angle)=-240)	OR OR OR OR OR OR
patch_rising_vs_setting ^d	Patch observed rising vs. setting			(w91_scan.hour_angle < 0)	
CES_gain_high_vs_low	CES average responsivity high vs. low			(SELECT AVG(timestream_gain_temp.fiducial_gain / diodes_gain_temp.gain_mean) FROM (w91_timestream AS timestream_gain_temp) JOIN (w91_diodes AS diodes_gain_temp) USING (module, diode) WHERE w91_timestream.run_id = timestream_gain_temp.run_id AND w91_timestream.run_subid = timestream_gain_temp.run_subid) > 0.951830) ^c	
diode_gain_high_vs_low	Diode average responsivity high vs. low			(w91_diodes.gain_mean 2.34083721782973)	>
sidelobe_el_high_vs_low	Sidelobe elevation high vs. low			(w91_scan.sidelobe_el > 35.475600) ^c	
array_orientation_90	Array orientation on the sky 90° apart			((w91_scan.array_orientation_equatorial >= -90. AND w91_scan.array_orientation_equatorial < 0.) OR (w91_scan.array_orientation_equatorial >= 90. AND w91_scan.array_orientation_equatorial < 180.))	OR OR OR OR

Continued on Next Page...

Table G.1 – Continued

Name	Description	Database Query
array_orientation_45	Array orientation on the sky 45° apart	$ \begin{aligned} & ((w91_scan.array_orientation_equatorial \\ & \geq -180.) \text{ AND } \\ & w91_scan.array_orientation_equatorial \\ & < -135.) \text{ OR } \\ & (w91_scan.array_orientation_equatorial \\ & \geq -90.) \text{ AND } \\ & w91_scan.array_orientation_equatorial \\ & < -45.) \text{ OR } \\ & (w91_scan.array_orientation_equatorial \\ & \geq 0.) \text{ AND } \\ & w91_scan.array_orientation_equatorial \\ & < 45.) \text{ OR } \\ & (w91_scan.array_orientation_equatorial \\ & \geq 90.) \text{ AND } \\ & w91_scan.array_orientation_equatorial \\ & < 135.) \end{aligned} $
fridge_cycle_commensurate	Scan period an integer multiple of refrigerator cycle period	$ \begin{aligned} & ((w91_scan.scan_period = 10) \text{ OR } \\ & (w91_scan.scan_period = 20) \text{ OR } \\ & (w91_scan.scan_period = 30) \text{ OR } \\ & (w91_scan.scan_period = 40)) \end{aligned} $
ambient_temperature_high_vs_low	Ambient temperature high vs. low	$(w91_atm.air_temp_med > -1.359500)^c$
sidelobe_sun ^e	Either sidelobe near sun vs. far	$ \begin{aligned} & (\text{IF}(w91_sidelobe_distance.sun_triple \\ & < w91_sidelobe_distance.sun_spillover, \\ & w91_sidelobe_distance.sun_triple, \\ & w91_sidelobe_distance.sun_spillover) \\ & \leq 55.827988)^c \end{aligned} $
sidelobe_moon ^e	Either sidelobe near moon vs. far	$ \begin{aligned} & (\text{IF}(w91_sidelobe_distance.moon_triple < \\ & w91_sidelobe_distance.moon_spillover, \\ & w91_sidelobe_distance.moon_triple, \\ & w91_sidelobe_distance.moon_spillover) \\ & \leq 69.572441)^c \end{aligned} $
sidelobe_galactic_latitude	Sidelobes near Galaxy vs. far	$ \begin{aligned} & (\text{IF}(\text{ABS}(w91_scan.sidelobe_lat) \\ & < \text{ABS}(w91_scan.sidelobe2_lat), \\ & \text{ABS}(w91_scan.sidelobe_lat), \\ & \text{ABS}(w91_scan.sidelobe2_lat)) \\ & \leq 10.656031)^c \end{aligned} $
scan_forward_and_backward	Left-going scans vs. right-going	N/A

Continued on Next Page...

Table G.1 – Continued

Name	Description	Database Query
scan_accel_and_decel	Accelerating vs. Decelerating	N/A

^a I excluded this test from normal analysis because we developed it after the null suite was finalized. I checked that the null suite passes even if I include these normally excluded tests.

^b I normally excluded these tests because they were designed to test the white noise correlation. Once I verified that the noise model was accurate, these tests did not target any suspected systematic error.

^c In order to make the divisions even, the thresholds are patch-dependent for some tests. For these, I give the thresholds for patch 2a.

^d This null test does not work for patch 7b because we only observed it while rising.

^e I excluded these tests from normal analysis because they were highly correlated with the Sun and Moon proximity tests.

To compute the correlation coefficient between two null tests, I used Eq. C.9 of [Bischoff \(2010\)](#).

Appendix H

List of Analysis Configurations

This appendix lists all analysis configurations I tried before settling on the final configuration¹. See §5 for discussion of the final configuration and how I chose it.

1. This very early configuration tested the software. I did not analyze the null-test results.
2. This configuration had no results because it used too much memory and processing crashed.
3. I reduced the number of null tests to 19 to avoid the processing crash. I found a negative-direction bias in the auto-correlation results indicating noise overestimation. The bias was strongest in the Q vs. U diode null test indicating a problem in white-noise correlation estimation.
4. This is the same as Configuration 3 except I used day-by-day cross-correlation instead of the usual pointing-based cross-correlation. It was an attempt to solve the bias, but it failed.
5. Due to a bug in previous configurations, I had not used the measured white-noise correlations. I fixed this bug, and the bias disappeared. I added the Q1U2 vs. Q2U2 and Q1U2 vs. Q2U1 null tests for future such noise-model problems. I updated the detector-angle calibration.
6. I updated the data selection (Buder, 2011d). To speed processing, I randomly used only 30% of the CESes. I removed two null tests (weather and cryostat temperature) because they were highly correlated with other existing tests. I updated the Type-B correction. I updated the cross-correlation divisions to 5 in azimuth and 8 in deck. I

1. I put more details in the report for each configuration at https://cmb.uchicago.edu/~ibuder/null_tests/null*/patch??/analysis/report.html

updated the beam calibration (Monsalve, 2011). The total $(\chi^{\text{null}})^2$ was too high (3% PTE). I found the auto-correlation bias at low ℓ in this configuration.

7. This is the same as 6 except I used all CESes.
8. I placed a tight (7 mHz) cut on knee frequency in attempt to solve the auto bias, but failed.
9. This is the same as 6 except I updated the data selection (Buder, 2011e), increased the number of bins for the ground filter from 20 to 40, updated the relative responsivity model (Ishidoshiro, 2011), and included 1/f-noise correlations within each module. There was a bug in the pointing model causing the collimation offset to be applied twice.
10. This is the same as 9 except I used all CESes. Patch 6a failed the null suite with 0% total $(\chi^{\text{null}})^2$ PTE .
11. I reverted the pointing model to 7; however, I did not analyze the results.
12. The same as 10 except I reduced the number of ground-filter bins to 20 in a failed attempt to solve the patch 6a failure.
13. I reverted the responsivity model to 7; however, I did not analyze the results.
14. I reverted the data selection to 7; however, it did not solve the patch 6a failure.
15. The same as 10 except I reverted both the data selection and pointing model to 7. It solved the patch 6a failure.
16. The same as 10 except I reverted both the data selection and responsivity model to 7. The patch 6a failure remained.

17. The same as 10 except I fixed the collimation-offset bug. This fixed the patch 6a failure.
18. The same as 10 except I updated the array pointing (Kusaka, 2011a). All patches passed the null suite.
19. I loosened the cuts on noise model χ^2 (13), allowing 3% more CES-diodes to pass.
20. The same as 18 except I used day-by-day cross-correlation. The lack of χ bias (similar to what we found in Q band) supported our decision to use simple Q–W cross-correlation in the combined analysis.
21. The same as 19 except I loosened the level-shift cut (16), allowing 1% more CES-diodes to pass.
22. The same as 19 except I loosened the cut on noise-power-spectrum spikes (17).
23. I tightened all cuts by an additional 13% of the CES-diodes.
24. The same as 18 except I used a frequency-domain high-pass filter instead of the slope subtraction filter.
25. The same as 24 except I randomly used only 30% of the CESes.
26. The same as 18 except I tightened the knee frequency (8), slopes (12), noise model (13), and level-shift (16) cuts. I removed an additional 8% of the CES-diodes.
27. The same as 18 except I tightened the slopes (12) and noise-model (13) cuts, rejecting an additional 4% of the CES-diodes. I updated the Type-B database (w91_typeb, not correction parameters) to reflect new CES definitions.
28. The same as 18 except I tightened the noise-model cut (13, below 1 Hz).
29. The same as 18 except I tightened the noise-model cut (14).

30. The same as 18 except I tightened the cut on noise-power-spectrum spikes (17).
31. The same as 18 except I tightened the cut on noise-power-spectrum spikes near 1.2 Hz (17).
32. The same as 18 except I loosened the noise-model cut (13, below 1 Hz).
33. The same as 18 except I loosened the noise-model cut (14).
34. The same as 30 except I tightened the cut further to 4.1σ .
35. The same as 31 except I tightened the cut further to 2.7σ .
36. The same as 34 except I tightened the cut further to 3.9σ .
37. The same as 18 except I tightened the 1-s level-shift cut (16).
38. The same as 18 except I loosened the 1-s level-shift cut (16). We also updated w91_timestream.
39. The same as 18 except I tightened the 10-s level-shift cut (16).
40. The same as 18 except I tightened the 100-s level-shift cut (16).
41. The same as 18 except I tightened the knee frequency cut (8).
42. The same as 18 except I averaged the two 50-Hz phase-switch states instead of differencing them. For ideal modules the result would have no signal and the same noise level as DD data. However, because of the phase-switch imbalance, there is worse 1/f noise than the DD data. The total $(\chi^{\text{null}})^2$ was too large.
43. The same as 18 except I tightened the slopes cut (12).
44. The same as 18 except I loosened the 10-s level-shift cut (16).

45. The same as 40 except I tightened the threshold to 18σ .
46. The same as 18 except I tightened the non-stationary noise cut (18).
47. The same as 18 except I increased the number of bins for the ground filter to 80 to test for residual ground pickup.
48. The same as 41 except I tightened the threshold to 0.9 times the scan frequency.
49. The same as 18 except I loosened the slopes cut.
50. The same as 45 except I tightened the threshold to 17σ .
51. The same as 18 except I loosened the non-stationary noise cut (18) threshold to 3.5.
52. The same as 18 except I updated the database (quiet_quality), array pointing, and Type-B correction to the final values.
53. The same as 18 except I updated the database (quiet_quality) to the final values.
54. The same as 52 except I tightened the noise-model, level-shift, power-spectrum spike, knee-frequency, and non-stationary noise cuts. The largest $(\chi^{\text{null}})^2$ PTE for patch 4a was 3%.
55. The same as 52 except I added the cut on scan-synchronous signal (19).
56. The same as 55 except I tightened the knee-frequency threshold to the scan frequency.
57. The same as 52 except I masked sources from the *WMAP* and *Planck* catalogs.
58. This was a test for new pipeline features. There were no null-test results.
59. The same as 55 except I updated the beam shape, absolute responsivity, and pointing mount model. I adjusted the null-test division thresholds to evenly divide the CES-diodes. The patch 4a total $(\chi^{\text{null}})^2$ was too large.

60. The same as 59 except I reverted the pointing model to 55. This fixed the patch 4a failure.

61. The same as 59 except I reverted the null-test thresholds to 55.

62. The same as 59 except I reverted the data selection to 52.

63. The same as 59 except I fixed a bug in the pointing model. The patch 4a total $(\chi^{\text{null}})^2$ for EE only was unexpectedly large.

64. The same as 59 except I reverted the data selection to 23.

65. The same as 59 except I removed deck flexure from the pointing model.

66. The same as 59 except I removed deck flexure and variable encoder offsets from the pointing model.

67. The same as 63 except I increased the number of ground filter bins to 80.

68. The same as 63 except I reverted the data selection to 23 and included the scan-synchronous signal cut.

69. The same as 63 except I reverted the data selection to 26 and included the scan-synchronous signal cut.

70. The same as 63 except I reverted the data selection to 19 and included the scan-synchronous signal cut.

71. The same as 63 except I doubled the number of azimuth divisions for cross-correlation (to 10).

72. The same as 63 except I changed the azimuth divisions for cross-correlation to separate CESes with different scan periods.

73. The same as 63 except I reduced the number of deck divisions for cross-correlation (to 4), pairing divisions 180° apart.
74. The same as 63 except I updated the beam shape and absolute responsivity. This was the final analysis configuration.

Appendix I

Statistical Uncertainty Scaling Between Configurations

When comparing two configurations with different data selection, the statistical uncertainty on the difference is smaller than the uncertainty for either configuration. To reduce the need for additional simulations, I developed a scaling method to estimate the expected statistical uncertainty. This method has two parts. First, I derived an approximate uncertainty for the case of instrument noise only (Buder, 2010a). Second, I extended the method to account for sample variance (Buder, 2012c).

I.1 Configuration-difference Error Bar

Let the two configurations be denoted 1 and 2. Let a denote the data common to 1 and 2; b, the data only accepted in 1; and c, the data only accepted in 2. Let x_1 be the power spectrum in Configuration 1, and x_a , the power spectrum computed from data a alone. Let σ_1 and σ_a be the corresponding statistical errors.¹ I assumed

$$x_1 = \frac{x_a \sigma_b^2 + x_b \sigma_a^2}{\sigma_a^2 + \sigma_b^2} \quad (\text{I.1})$$

as a heuristic approximation to the weighting of data a and b in Configuration 1. By symmetry,

$$x_2 = \frac{x_a \sigma_c^2 + x_c \sigma_a^2}{\sigma_a^2 + \sigma_c^2}. \quad (\text{I.2})$$

The power spectrum difference

$$x_1 - x_2 = \frac{(x_a \sigma_b^2 + x_b \sigma_a^2)(\sigma_a^2 + \sigma_c^2) - (x_a \sigma_c^2 + x_c \sigma_a^2)(\sigma_a^2 + \sigma_b^2)}{W}, \quad (\text{I.3})$$

1. In this case the instrument-noise-only error from noise simulations should be used.

where

$$W \equiv (\sigma_a^2 + \sigma_b^2)(\sigma_a^2 + \sigma_c^2). \quad (\text{I.4})$$

Simplifying,

$$x_1 - x_2 = \frac{x_a \sigma_a^2 \sigma_b^2 + x_b \sigma_a^4 + x_b \sigma_a^2 \sigma_c^2 - x_a \sigma_a^2 \sigma_c^2 - x_c \sigma_a^4 - x_c \sigma_a^2 \sigma_b^2}{W}. \quad (\text{I.5})$$

Since a, b, and c are disjoint, they have independent fluctuations. Therefore their error bars on the difference add in quadrature as

$$\sigma_{x_1 - x_2} = \sqrt{A^2 + B^2 + C^2}, \quad (\text{I.6})$$

where

$$A \equiv \frac{\sigma_a^3 (\sigma_b^2 - \sigma_c^2)}{W} \quad (\text{I.7})$$

$$B \equiv \frac{\sigma_b (\sigma_a^4 + \sigma_a^2 \sigma_c^2)}{W} \quad (\text{I.8})$$

$$C \equiv \frac{\sigma_c (\sigma_a^4 + \sigma_a^2 \sigma_b^2)}{W}. \quad (\text{I.9})$$

B and C can be attributed to the difference in data between 1 and 2. A is the fluctuation of data a times the change in its weighting between 1 and 2. Since I usually investigated small changes in data selection,

$$\sigma_b, \sigma_c \gg \sigma_a \quad (\text{I.10})$$

so typically $B, C \gg A$.

Usually I computed σ_1 and σ_2 are from noise simulations. To be consistent with the weighting used in Eq. I.1 I assumed

$$\frac{1}{\sigma_1^2} = \frac{1}{\sigma_a^2} + \frac{1}{\sigma_b^2}. \quad (\text{I.11})$$

I further assumed that the error bar for different data selections was inversely proportional to the data size

$$\frac{\sigma_1}{\sigma_a} = \frac{D_a}{D_a + D_b}, \quad (\text{I.12})$$

where D_a is the number of CES-diodes in data a. I used Eq. I.12 to determine σ_a when σ_1, D_a, D_b are known². Combining Eqs. I.11 and I.12 I solved for

$$\sigma_b = \sigma_1 \frac{D_a + D_b}{\sqrt{2D_a D_b + D_b^2}}. \quad (\text{I.13})$$

The corresponding equation for

$$\sigma_c = \sigma_2 \frac{D_a + D_c}{\sqrt{2D_a D_c + D_c^2}} \quad (\text{I.14})$$

is found by replacing $1 \rightarrow 2$ and $b \rightarrow c$.

I.2 Impact of CMB–instrument-noise Accidental Correlation

Above I assumed that sample variance (the impact of non-zero CMB power) was negligible. However, that is not strictly true. In the configuration comparison tests, I differenced two power spectra computed from the same patch. Both configurations have the same CMB realization. However, the instrument noise is different, and there is accidental correlation between the noise and CMB signal. Although this correlation was 0 on average, the accidental correlation increased the statistical fluctuation of the difference between the spectra from different configurations³. First I calculated this increased fluctuation for a toy model based

2. One could instead replace $1 \rightarrow 2, b \rightarrow c$ for a different estimate of σ_a . The difference between them measures the error introduced by the heuristic assumptions about σ . I checked this difference was $\approx 10\%$ or less.

3. This effect would not have arisen if I had differenced the maps of two configurations and computed the (null) spectra of the difference map. For historical reasons, the pipeline did not have this facility until very late in the analysis. For simplicity, I conducted all configuration difference tests by differencing power

on a single spatial mode on the sky. Then I extended the model for the case of multiple modes. Finally, I provided additional details for the EB spectrum.

I.2.1 CMB–instrument-noise Correlation Toy Model

Suppose there is only a single mode a on the sky. We measure it with some noise

$$\hat{a} = a^{\text{CMB}} + a^N. \quad (\text{I.15})$$

The measured spectrum is

$$\hat{C} = \hat{a}^2 = (a^{\text{CMB}})^2 + (a^N)^2 + 2a^{\text{CMB}}a^N. \quad (\text{I.16})$$

The fluctuation of $(a^{\text{CMB}})^2$ is the sample variance, which I ignore because it is the same for two different configurations. The fluctuation of $(a^N)^2$ is the noise variance, which I treated in Appendix I.1. (Our cross-correlation, §5.4.3, eliminates the average value of the noise variance—the noise bias—from the result, but we still have the fluctuation.) The fluctuation of $2a^{\text{CMB}}a^N$ is the new variance I will solve for. Since we have a fixed sky realization, I treat a^{CMB} as a constant. Thus the unknown is the expected variation Δa^N . The heuristic error bar gives $\Delta^{\text{cut}}C^N$, the fluctuation of the noise bias between two configurations. Now I relate it to $\Delta^{\text{cut}}a^N$.

$$C^N = (a^N)^2 \quad (\text{I.17})$$

$$\Delta C^N = 2a^N \Delta a^N \quad (\text{I.18})$$

$$\Delta a^N = \frac{\Delta C^N}{2\sqrt{C^N}} \quad (\text{I.19})$$

$$\Delta^{\text{cut}}a^N = \frac{\Delta^{\text{cut}}C^N}{2\sqrt{C^N}} \quad (\text{I.20})$$

spectra. I suggest that future experiments should conduct this type of test by differencing maps.

The fluctuation of $2a^{\text{CMB}}a^N$ is

$$\Delta 2a^{\text{CMB}}a^N = 2a^{\text{CMB}}\Delta a^N = 2a^{\text{CMB}}\frac{\Delta C^N}{2\sqrt{C^N}} = \sqrt{\frac{C^{\text{CMB}}}{C^N}}\Delta C^N, \quad (\text{I.21})$$

where C^{CMB} is the (true) CMB power, C^N is the noise bias, and ΔC^N is the noise error bar due to the configuration difference. Both ΔC^N and this new fluctuation are proportional to Δa^N ; however, the constants of proportionality (a^N and a^{CMB}) are uncorrelated. Therefore, when averaging over many modes, the two fluctuations add in quadrature. The total fluctuation is $\sqrt{1 + \frac{C^{\text{CMB}}}{C^N}}\Delta C^N$.

I.2.2 Multiple Sky Modes

Now I consider the more general case of a bandpower C_b which averages n_b modes. I assume the modes have identical statistical properties. (In practice they would have a range of ℓ and unequal weighting.) The measured spectrum becomes

$$\hat{C}_b = \langle \hat{a}_{\ell m}^2 \rangle = \langle (a^{\text{CMB}})^2 + (a^N)^2 + 2a^{\text{CMB}}a^N \rangle, \quad (\text{I.22})$$

where $\langle \rangle$ indicates the average over n_b modes. I modified Eq. I.18 to

$$\Delta C_b^N = 2a_{\ell m}^N \Delta a_{\ell m}^N / \sqrt{n_b} \quad (\text{I.23})$$

because the n_b modes are independent so their fluctuations add in quadrature. (Equivalently, C_b^N is an average of n_b measurements so the fluctuation of the average $\sim 1/\sqrt{n_b}$.) Hence Eq. I.20 becomes

$$\Delta^{\text{cut}} a_{\ell m}^N = \frac{\sqrt{n_b} \Delta^{\text{cut}} C_b^N}{2\sqrt{C_b^N}}. \quad (\text{I.24})$$

The additional fluctuation is now

$$\begin{aligned}\Delta \langle 2a_{\ell m}^{\text{CMB}} a_{\ell m}^N \rangle &= \frac{1}{\sqrt{n_b}} \Delta 2a_{\ell m}^{\text{CMB}} a_{\ell m}^N = \frac{1}{\sqrt{n_b}} 2a_{\ell m}^{\text{CMB}} \Delta a_{\ell m}^N = \\ \frac{1}{\sqrt{n_b}} 2a_{\ell m}^{\text{CMB}} \frac{\sqrt{n_b} \Delta^{\text{cut}} C_b^N}{2\sqrt{C_b^N}} &= \sqrt{\frac{C_b^{\text{CMB}}}{C_b^N}} \Delta^{\text{cut}} C_b^N,\end{aligned}\quad (\text{I.25})$$

which is equivalent to Eq. I.21.

I.2.3 EB Power Spectrum

The EB spectrum

$$\hat{C}_{EB} = \hat{a}_E \hat{a}_B = (a_E^N + a_E^{\text{CMB}}) a_B^N = a_E^N a_B^N + a_E^{\text{CMB}} a_B^N \quad (\text{I.26})$$

is complicated because while there is no signal (I assume the CMB B mode is negligible compared to our noise level), the accidental correlation between the CMB E mode (a_E^{CMB}) and the noise B mode (a_B^N) increases the variance. In the instrument-noise-only case, the error bar is

$$\Delta^{\text{noise}} C_{EB} = \sqrt{(a_B^N \Delta a_E^N)^2 + (a_E^N \Delta a_B^N)^2} \quad (\text{I.27})$$

because a_E^N and a_B^N are independent so the two terms add in quadrature.

The error due to the accidental correlation is

$$\Delta a_E^{\text{CMB}} a_B^N = a_E^{\text{CMB}} \Delta a_B^N = \frac{\sqrt{C_{EE}^{\text{CMB}}} \Delta C_{BB}^N}{2\sqrt{C_{BB}^N}}, \quad (\text{I.28})$$

where ΔC_{BB}^N is the BB spectrum error bar (e.g. due to the cut change) and I used Eq. I.20 to rewrite Δa_B^N . This error is uncorrelated with $\Delta^{\text{noise}} C_{EB}$ for the reason given in

Appendix I.2.1. The total fluctuation is therefore

$$\sqrt{(\Delta^{\text{noise}} C_{EB})^2 + \left(\frac{\sqrt{C_{EE}^{\text{CMB}}} \Delta C_{BB}^N}{2\sqrt{C_{BB}^N}} \right)^2}.$$

Appendix J

Type-B-glitch—correction-parameter File Format

This sample Type-B-glitch—correction-parameter file demonstrates the parameter-file format for the Type-B-glitch—correction software.

```
--FORMAT_VERSION2--  
#Required at the beginning, sets version of the parameter format  
  
#Any line beginning with '#' is a comment  
  
# Sampling rate (Hz) to be used for noise calculation  
--DEMOD_SAMPLING_RATE-- 100  
  
# Iteration count for Newton's method (solving Delta G)  
--N_ITERATIONS-- 3  
  
# Factor to be multiplied to the heights below  
--HEIGHT_FACTOR-- 2.0  
  
# Multi-glitch correction. How many noise sigma from the glitch to correct.  
--MULTI_GLITCH_COVERAGE_SIGMA-- 5.0  
  
# Preamp low-pass filter RMS reduction "Preamp factor"  
--PREAMP_RMS_FACTOR--  
  
# Module Q1 U1 U2 Q2  
  
0 0.644467 0.643774 0.637945 0.632578  
1 0.576704 0.572848 0.741313 0.607144
```

```

#Modules cannot be skipped

# Specify center-height pattern. -1 is non-glitching.

--CENTER_HEIGHT_PATTERN--

# Module Q1 U1 U2 Q2

00      0      0      0      0
01      -1     0      0      1

#Modules cannot be skipped

# Detector RMS (bit*sqrt[s]) to be used as default

# Not used in iterative correction

--DETECTOR_RMS--

#      Module   Q1      U1      U2      Q2
00      0.20    0.20    0.20    0.20
01      0.20    0.20    0.20    0.20

#Modules cannot be skipped

# Center-Height data

# Normal pattern

--CENTERS_HEIGHTS--      0

#Location (bits, 0 is the center of the ADC range)

#                           error(ignored)
#                           height(bits)
#                           error(ignored)

-65500.300      0.0      11.732002    0.0

```

-64455.700	0.0	-6.8432889	0.0
-63452.300	0.0	11.732002	0.0
-62407.700	0.0	-0.97366601	0.0

#Location must increase monotonically

Pattern for other diodes

__CENTERS_HEIGHTS__	1		
-65501.300	0.0	11.732002	0.0
-64455.700	0.0	-6.8432889	0.0
-63452.300	0.0	11.732002	0.0
-62407.700	0.0	-1.97366601	0.0

#More patterns

References

2012, inb(2) - Linux man page, <http://linux.die.net/man/2/inb> 63

2012, ioperm(2) - Linux man page, <http://linux.die.net/man/2/ioperm> 63

Aumont, J., Conversi, L., Thum, C., et al. 2010, *A&A*, 514, A70 93

Bannister, P., Nixon, G., & Staggs, S. 2006, Final Preproduction Memo for QUIET W-Band OMTs, <https://cmb.uchicago.edu/project/images/f/fa/FinalWOMT-compressed.pdf>, QUIET internal memo 20

Barkats, D., Bischoff, C., Farese, P., et al. 2005, *ApJS*, 159, 1 22

Battye, R. A., Browne, I. W. A., Peel, M. W., Jackson, N. J., & Dickinson, C. 2011, *MNRAS*, 413, 132 152

Baumann, D., Jackson, M. G., Adshead, P., et al. 2009, in *American Institute of Physics Conference Series*, Vol. 1141, *American Institute of Physics Conference Series*, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith, 10 6

Becker, M. 2008, Optimization of quiet polarimeters, QUIET internal memo 23

Bischoff, C. 2006a, Addendum to Preamp Interface Spec v2.0, <https://cmb.uchicago.edu/project/images/8/88/Preamp-2.0-addendum.ps>, QUIET internal memo 44

—. 2006b, Switched signal from a biased detector diode, https://cmb.uchicago.edu/project/images/f/f6/Preamp_signalcalculation.ps, QUIET internal memo 25, 37

—. 2007, Measurements of the First Item W-band MAB, <https://cmb.uchicago.edu/project/images/6/6e/Mabtest.ps>, QUIET internal memo 31

—. 2009, Chicago Pointing and Coordinate Conventions v2.1, https://cmb.uchicago.edu/project/images/a/ad/Chicago_pointing-2.ps, QUIET internal memo 96

Bischoff, C., & Kapner, D. 2006, W-band MAB Version 2 Specification, <https://cmb.uchicago.edu/project/images/e/e9/MABspecs-2.pdf>, QUIET internal memo 31

Bischoff, C., & Kusaka, A. 2009, Derivation of White Noise Correlation in QUIET Detector Diodes, https://cmb.uchicago.edu/project/images/a/a8/Diode_noise_correlation.pdf, QUIET internal memo 25, 102, 206

Bischoff, C., Hyatt, L., McMahon, J. J., et al. 2008, *ApJ*, 684, 771 12

Bischoff, C. A. 2010, PhD thesis, The University of Chicago, http://quiet.uchicago.edu/depot/pdf/bischoff_thesis.pdf 13, 15, 25, 32, 37, 41, 44, 53, 212

Bogdan, M. 2005, The QUIET Electronics Box Implementation - I, https://cmb.uchicago.edu/project/images/1/1c/Design_of_QEB_1.pdf, QUIET internal memo 39

Bornemann, J., & Labay, V. 1995, Microwave Theory and Techniques, IEEE Transactions on, 43, 1782 20

Boyle, L. A., Steinhardt, P. J., & Turok, N. 2006, Phys. Rev. Lett., 96, 111301 6

Brizius, A. 2009, Type-B Glitch Correction, https://cmb.uchicago.edu/project/images/1/1b/Type-B_Glitch_Fitting.pdf, QUIET internal memo 99

—. 2010, Observing Priorities, https://cmb.uchicago.edu/project/index.php/Observing_Priorities, QUIET internal memo 68

Brizius, A. R. 2011, PhD thesis, The University of Chicago, <http://search.proquest.com/docview/894759190?accountid=14657> 13, 18, 22, 23, 26, 31, 32, 36, 41, 49, 81

Brown, M. L., Ade, P., Bock, J., et al. 2009, ApJ, 705, 978 8, 12

Buder, I. 2008, QUIET Electronics Enclosure Temperature Control System, <https://cmb.uchicago.edu/project/images/4/4b/PIDControl.pdf>, QUIET internal memo 41

—. 2010a, Heuristic Estimate of the Statistical Fluctuation of the Power Spectrum Difference Between Different Cuts, https://cmb.uchicago.edu/project/images/8/83/Difference_error_bar_heuristic-2010-07-19.pdf, QUIET internal memo 220

—. 2010b, Patch Consistency χ^2 Test, https://cmb.uchicago.edu/project/images/3/32/Patch_consistency-2010-11-19.pdf, QUIET internal memo 136

—. 2010c, Q19 Array Sensitivity, https://cmb.uchicago.edu/project/images/4/4b/Q19_array_sensitivity.pdf, QUIET internal memo 205

—. 2010d, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V, Proc. SPIE, 7741, 77411D 12

—. 2011a, 800-kHz Non-Gaussianity Quick Study, <https://cmb.uchicago.edu/project/images/a/a9/Gaussianity-2011-04-11.pdf>, QUIET internal memo 148

—. 2011b, Low ℓ Auto Bias in Null Suite, https://cmb.uchicago.edu/project/images/f/fc/Auto_bias-2011-06-20.pdf, QUIET internal memo 134

—. 2011c, Type-B Systematic Error for W Band, https://cmb.uchicago.edu/project/images/6/6b/Typeb_syst_w-2011-11-08.pdf, QUIET internal memo 143

—. 2011d, W-band Data Viewer/Data Quality Study, https://cmb.uchicago.edu/~ibuder/w91_viewer_study/study2.html, QUIET internal memo 115, 213

—. 2011e, W-band Data Viewer/Data Quality Study 2, https://cmb.uchicago.edu/~ibuder/w91_viewer_study/study3.html, QUIET internal memo 115, 214

- . 2011f, W-band Reference Period Viewer Study, https://cmb.uchicago.edu/~ibuder/w91_viewer_study/study.html, QUIET internal memo 115
- . 2011g, W-Band Type-B Correction v4, <https://cmb.uchicago.edu/project/images/5/59/Typeb-2011-03-28.pdf>, QUIET internal memo 100
- . 2012a, Configuration Comparison, 52 vs. 57, patch7b, https://cmb.uchicago.edu/~ibuder/null_tests/config_difference_52_57/patch7b/config_comparison.html, QUIET internal memo 152
- . 2012b, Configuration Comparison, 63 vs. 67, patch4a, https://cmb.uchicago.edu/~ibuder/null_tests/config_difference_63_67/patch4a/config_comparison.html, QUIET internal memo 151
- . 2012c, Fluctuation of CMB–Instrument-Noise Correlation, https://cmb.uchicago.edu/project/images/0/08/Cmb_instrument_noise_correlation.pdf, QUIET internal memo 220

Buder, I., & Kusaka, A. 2011, 1/f Noise Correlation and Modeling, https://cmb.uchicago.edu/project/images/f/fe/1f_correlation-2011-08-02.pdf, QUIET internal memo 101

Buder, I., et al. 2011, Strategy for Null-Test Selection and Evaluation, https://cmb.uchicago.edu/project/images/a/ad/Null_suite_strategy-2011-05-09.pdf, QUIET internal memo 128

- . 2012, Null-Test Report for Configuration 74 (patch 2a), <https://cmb.uchicago.edu/~ibuder/null74/patch2a/analysis/report.html>, QUIET internal memo 208

Chiang, H. C., Ade, P. A. R., Barkats, D., et al. 2010, *ApJ*, 711, 1123 8, 12

Chicago. 2010, Shutoff/Turnon instructions for the W-band Receiver, https://cmb.uchicago.edu/site_document/trunk/ReceiverOnOff_W.pdf, QUIET internal memo 68

Chinone, Y. 2011a, PhD thesis, Tohoku University, http://quiet.uchicago.edu/depot/pdf/chinone-2011_thesis.pdf 13, 91

- . 2011b, W-band Far-Sidelobe Analysis, https://cmb.uchicago.edu/project/images/1/14/W-band_far-sidelobe_document_2011-11-08_YujiChinone.pdf, QUIET internal memo 20, 115, 149

Chinone, Y., Nagai, M., & Tajima, O. 2011, Detector Angle for W-band Analysis, https://cmb.uchicago.edu/project/images/8/8a/W91_angle_pcl.pdf, QUIET internal memo 93

Chinone, Y., Radford, S., Jara, C., et al. 2012, W-band Bandpass, https://cmb.uchicago.edu/project/images/1/19/W91_bandpass_20110208.pdf, QUIET internal memo 81

Cleary, K. A. 2010, Proc. SPIE, 7741, 77412H 22

Cooray, A., Melchiorri, A., & Silk, J. 2003, Physics Letters B, 554, 1 11

Croft, R. A. C., & Dailey, M. 2011, ArXiv e-prints, arXiv:1112.3108 [astro-ph.CO] 127

CyberPower. 2012, CyberSwitch Network Management System, http://www.cyberpowersystems.com/user-manuals/CSW8RU_UM.pdf, Product Manual Beta3_6-16 64

Delabrouille, J., et al. 2012, *Astronomy & Astrophysics*, to appear 152

Dodelson, S. 2003, AIP Conf. Proc., 689, 184 8

Dumoulin, R. 2007, AIB Cable Proposal, <https://cmb.uchicago.edu/project/images/2/23/AIBCABLES2.pdf>, QUIET internal memo 32

Dumoulin, R. N. 2011, PhD thesis, Columbia University, <http://hdl.handle.net/10022/AC:P:11789> 177

Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52 8

Dunkley, J., Amblard, A., Baccigalupi, C., et al. 2009, in American Institute of Physics Conference Series, Vol. 1141, American Institute of Physics Conference Series, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith, 222 12

Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, ApJ, 426, 23 123

Fixsen, D. J. 2009, ApJ, 707, 916 7

Fixsen, D. J., Cheng, E. S., Gales, J. M., et al. 1996, ApJ, 473, 576 7

Gaier, T., Lawrence, C. R., Seiffert, M. D., et al. 2003, New Astronomy Reviews, 47, 1167, proceedings of the Workshop on The Cosmic Microwave Background Radiation and its Polarization 22

Gaier, T. C. 2003, Polarimetry in Astronomy, Proc. SPIE, 4843, 296 15

GE Fanuc Embedded Systems. 2007, VMIVME-7700, <https://cmb.uchicago.edu/project/images/5/52/VME7700.pdf>, product manual, Rev. F 53

Gold, B., Odegard, N., Weiland, J., et al. 2011, ApJS, 192, 15 95, 152

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759 xii, 121

Gundersen, J., & Wollack, E. 2009, Journal of Physics: Conference Series, 155, 012005 20

Güsten, R., Nyman, L. Å., Schilke, P., et al. 2006, A&A, 454, L13 73

Guth, A. H. 1981, *Phys. Rev. D*, **23**, 347 2

Hamimeche, S., & Lewis, A. 2008, *Phys. Rev. D*, **77**, 103013 153, 154

Hansen, F. K., & Górska, K. M. 2003, *MNRAS*, **343**, 559 121

Hasegawa, M. 2011a, Bad ces-diode rejection on preamp power supply problem, <https://cmb.uchicago.edu/project/images/4/4c/Main.pdf>, QUIET internal memo 115

—. 2011b, Dataset for W-band Analysis, https://cmb.uchicago.edu/project/images/9/94/W91_dataset_doc.pdf, QUIET internal memo 75

—. 2011c, Observation summary, https://cmb.uchicago.edu/project/images/1/18/QUIET_Wband_Observation_Summary_111025.pdf, QUIET internal memo 75

—. 2011d, Observation Summary, https://qufs.uchicago.edu/DQM/live_Wband/trunk/html/daily_log/observation_summary_based_on_live_encoder.html, QUIET internal memo 75

Hedman, M. M., Barkats, D., Gundersen, J. O., Staggs, S. T., & Winstein, B. 2001, *ApJ*, **548**, L111 22

Hivon, E., Górska, K. M., Netterfield, C. B., et al. 2002, *ApJ*, **567**, 2 121

Hu, W., & Dodelson, S. 2002, *Ann. Rev. Astron. Astrophys.*, **40**, 171 7, 9

Hu, W., & White, M. 1997, *New Astronomy*, **2**, 323 9

IEEE and The Open Group. 2012a, `pthread_cond_wait(3)` - Linux man page, http://linux.die.net/man/3/pthread_cond_wait 174

—. 2012b, `pthread_mutex_lock(3)` - Linux man page, http://linux.die.net/man/3/pthread_mutex_lock 58

Imbriale, W., Gundersen, J., & Thompson, K. 2011, *Antennas and Propagation, IEEE Transactions on*, **59**, 1972 18

Ishidoshiro, K. 2011, Modeling of Relative Responsivities in PCL pipeline – Application for W-band data set, <https://cmb.uchicago.edu/project/images/b/b8/Relativégain20110802.pdf>, QUIET internal memo 92, 214

—. 2012, Validation of Statistical Errors by using Random Nulls, https://cmb.uchicago.edu/project/images/2/26/Error_estimation_2012_0306.pdf, QUIET internal memo 134

Jarosik, N., Bennett, C. L., Halpern, M., et al. 2003, *ApJS*, **145**, 413 22

Kapner, D. 2006, Understanding the QUIET Detector Diodes, <https://cmb.uchicago.edu/project/images/1/1d/Diode.pdf>, QUIET internal memo 25

Keisler, R., Reichardt, C., Aird, K., et al. 2011, *ApJ*, 743, 28 8

Klein, J. R., & Roodman, A. 2005, *Annual Review of Nuclear and Particle Science*, 55, 141 127

Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, *ApJS*, 192, 18 5, 6, 7, 8, 155, 160

Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, *Nature*, 420, 772 12

Kusaka, A. 2007, Leakage Effects due to Receiver Imperfections (Supplementation for the Appendix of Alex's Thesis), https://cmb.uchicago.edu/project/images/9/92/Alex_appendix_suppl.pdf, QUIET internal memo 25

—. 2008, Simulation Study of the Systematic Effects Due to TOD Filter in Pseudo- C_ℓ Analysis, https://cmb.uchicago.edu/project/images/6/6e/Pseudo_cl_tod_systematics.pdf, QUIET internal memo 125

—. 2009, Map Making Calculations Using Four Diodes, QUIET internal memo 121

—. 2010a, Power Spectrum Calculation Using “Each-Az-Filtered” TOD, https://cmb.uchicago.edu/project/images/4/4a/Each_az_filter_spectrum-v1.pdf, QUIET internal memo 116

—. 2010b, Timing Offset Monitor, https://cmb.uchicago.edu/project/index.php/Timing_Offset_Monitor, QUIET internal memo 59

—. 2010c, Type-B Glitch Correction, https://cmb.uchicago.edu/project/images/9/98/Adc_anomaly-v3.pdf, QUIET internal memo 97, 99, 101, 146

—. 2011a, Array Pointing Using Moon, https://cmb.uchicago.edu/project/images/c/cd/Array_pointing_summary-2011-09-27.pdf, QUIET internal memo 215

—. 2011b, Moon analysis for array pointing Summary, https://cmb.uchicago.edu/project/images/e/e9/Summary_of_moon_analysis_for_array_pointing.pdf, QUIET internal memo 96

—. 2011c, W-band Weather Statistics, https://cmb.uchicago.edu/project/images/5/55/Wband_weather_stat.pdf, QUIET internal memo 116

—. 2012a, moon_pointing.pdf, https://cmb.uchicago.edu/project/images/1/14/Moon_pointing.pdf, QUIET internal memo 96

—. 2012b, Upsampled MC summary, https://cmb.uchicago.edu/project/images/b/be/Upsampled_mc-summary.pdf, QUIET internal memo 125

Kusaka, A., et al. 2012, W-band Noise Property, https://cmb.uchicago.edu/project/images/1/1b/Wband_noise_property_summary-20120302.pdf, QUIET internal memo 102, 205, 207

Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, *ApJS*, 192, 16 8

Leitch, E., et al. 2002, *Nature*, 420, 763 22

Liddle, A. R., & Lyth, D. H. 1993, *Physics Reports*, 231, 1 6

Linde, A. D. 2005, *arXiv:hep-th/0503203 [hep-th]* 3

—. 2008, *Lect.Notes Phys.*, 738, 1 2

Mizugutchi, Y., Akagawa, M., & Yokoi, H. 1976, in IEEE Antennas and Propagation Society International Symposium, 1976, Vol. 14, 2 18

Monsalve, R. 2011, W-Band Gain and Beam Analysis with Tau A, https://cmb.uchicago.edu/project/images/b/b7/Miami_taua_report_051711.pdf, QUIET internal memo 214

—. 2012a, PhD thesis, University of Miami 89, 122

—. 2012b, W-Band Beams, https://cmb.uchicago.edu/project/images/7/73/Miami_wband_report.pdf, QUIET internal memo 122

Naess, S. 2012, PhD thesis, University of Oslo 88

Nagai, M. 2011a, Gain Model for W-band Analysis, https://cmb.uchicago.edu/project/images/d/d4/GainModel_20110722.pdf, QUIET internal memo 91

—. 2011b, Model for Anchor Gain – Absolute Gain Model for W-band Analysis, <https://cmb.uchicago.edu/project/images/4/43/AbsoluteGain.pdf>, QUIET internal memo 91

—. 2012a, Modeling of Pol. SSS Az. Structure, https://cmb.uchicago.edu/project/images/4/47/Pol_SSS_20120214.pptx.pdf, QUIET internal memo 117, 150

—. 2012b, Systematic errors for detector angle in W band, https://cmb.uchicago.edu/project/images/c/c3/Systematics_angle_conf09_20120228.pdf, QUIET internal memo 141

—. 2012c, Systematic errors for ground pickup including variation, https://cmb.uchicago.edu/project/images/4/4e/Systematics_ground_pickup_variation_20120313.pdf, QUIET internal memo 151

—. 2012d, Systematic errors for pol. SSS variation, https://cmb.uchicago.edu/project/images/7/75/Pol_sss_variation_upper_limit_20120313.pdf, QUIET internal memo 150

—. 2012e, Systematic errors induced by detector gain uncertainties in W-band, https://cmb.uchicago.edu/project/images/8/89/Systematics_gain_20120205.pdf, QUIET internal memo 41, 92, 93, 141

Nagai, M., & Chinone, Y. 2011a, Input Parameters for W-band Receiver Calibration with Tau A, <https://cmb.uchicago.edu/project/images/8/84/TauAInputValues.pdf>, QUIET internal memo 91

—. 2011b, List of Tau A Data Sets in W-band Observations, <https://cmb.uchicago.edu/project/images/b/bb/TauADataSetList.pdf>, QUIET internal memo 77

Newburgh, L. 2010, PhD thesis, Columbia University, http://quiet.uchicago.edu/depot/pdf/newburgh_thesis.pdf 25

Newburgh, L., Miller, A., & Williamson, R. 2006, QUIET Phase I Cryostats: W and Q Band, Version 10.3, https://cmb.uchicago.edu/project/images/8/8a/Cryo_v10_3.pdf, QUIET internal memo 26

Nguyen, H. 2009, QUIET Shift Utility GUI Script Instructions, <http://home.fnal.gov/~7Ehogann/QUIET/Shift%20Utility%20Instructions.pdf>, QUIET internal memo 68

Nixon, G. 2006, FPC_Preamppinout.pdf, https://cmb.uchicago.edu/project/images/8/8e/FPC_Preamppinout.pdf, QUIET internal memo 32

—. 2007, AIB_pinout.xls, https://cmb.uchicago.edu/project/images/1/16/AIB_pinout.xls, QUIET internal memo 32

—. 2008, OMT measurements, https://cmb.uchicago.edu/project/images/e/e4/0mt_measurements_080308.pdf, QUIET internal memo 22

Nixon, G., & Dumoulin, R. 2006, FPC Redesign Proposals, https://cmb.uchicago.edu/project/images/9/9f/Fpc_revision_101106.pdf, QUIET internal memo 32

Oslo. 2012, The QUIET W-band ML mount model, https://cmb.uchicago.edu/project/images/f/f1/Pointing_Wband.pdf, QUIET internal memo 94

Padin, S., Shepherd, M. C., Cartwright, J. K., et al. 2002, *PASP*, 114, 83–18, 22

Pardo, J., Cernicharo, J., & Serabyn, E. 2001, *Antennas and Propagation, IEEE Transactions on*, 49, 1683–72

PCL group. 2010, Modeling the PDF of the Pseudo- C_ℓ Estimator, QUIET internal memo 154

Planck Collaboration. 2011, *A&A*, 536, A7–152

QUIET Chicago Group. 2008, 25ms_delay.pdf, https://cmb.uchicago.edu/project/images/1/12/25ms_delay.pdf, QUIET internal memo 59

QUIET Collaboration. 2005, moduledoc, <https://cmb.uchicago.edu/project/images/c/c8/Moduledoc.pdf>, QUIET internal memo 23, 25

- . 2010, Scheduling Manual, https://cmb.uchicago.edu/site_document/trunk/scheduling_manual.pdf, QUIET internal memo 68
- . 2012a, ArXiv e-prints, [arXiv:1207.5034 \[astro-ph.CO\]](https://arxiv.org/abs/1207.5034) 13
- . 2012b, ArXiv e-prints, [arXiv:1207.5562 \[astro-ph.IM\]](https://arxiv.org/abs/1207.5562) 12, 15, 17, 22, 37, 72, 149, 151
- QUIET Collaboration, Bischoff, C., Brizius, A., et al. 2011, *ApJ*, 741, 111 12, 13, 72, 96, 152
- Reichardt, C., Shaw, L., Zahn, O., et al. 2011, [arXiv:1111.0932 \[astro-ph.CO\]](https://arxiv.org/abs/1111.0932) 8
- Ryden, B. 2002, Introduction to Cosmology (Benjamin Cummings), <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0805389121> 1, 7
- Sachs, R. K., & Wolfe, A. M. 1967, *ApJ*, 147, 73 7
- Samtleben, D. 2008, Summary of patch selection, <https://cmb.uchicago.edu/project/images/0/0b/Reg1.pdf>, QUIET internal memo 74
- Sanders, J. R. 2010, Sidelobe Measurement and Reduction Efforts for QUIET Phase I, <https://cmb.uchicago.edu/project/images/e/ed/Sidelobememo.pdf>, QUIET internal memo 20
- Scott, D., & Smoot, G. 2006, [arXiv:astro-ph/0601307](https://arxiv.org/abs/astro-ph/0601307) 7
- Scott, D., Srednicki, M., & White, M. 1994, *ApJ*, 421, L5 8
- Seiffert, M. 2008, "QUIET MMIC Bias mod A0", https://cmb.uchicago.edu/project/images/9/9d/Mmic_ver3_prototype_schematic.pdf, QUIET internal memo 36
- Shepherd, M. 2007, Real-time motion-control modifications for QUIET, https://cmb.uchicago.edu/site_document/trunk/quiet_drive_system.pdf, QUIET internal memo 76
- . 2008a, The CBI Register List, <http://www.astro.caltech.edu/~mcs/Chajnantor/QUIET/help/registers.html> 64
- . 2008b, Commanding smooth telescope trajectories, https://cmb.uchicago.edu/site_document/trunk/trajectories.pdf, QUIET internal memo 76
- . 2008c, The QUIET telescope drive-system, http://www.astro.caltech.edu/~mcs/Chajnantor/QUIET/drive_system.pdf, QUIET internal memo 96
- Sievers, J. L., Achermann, C., Bond, J., et al. 2007, *ApJ*, 660, 976 12
- Smith, K. 2009, Null tests in pseudo- C_ℓ pipelines, <https://cmb.uchicago.edu/project/images/b/b3/09-04-28-nulltests.ps>, QUIET internal memo 129

Smith, K. M., Cooray, A., Das, S., et al. 2009, in American Institute of Physics Conference Series, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith, Vol. 1141, 121–12

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJ*, 396, L1–7

Symmetricom Inc. 2005, TTM635VME/TTM350VXI Time and Frequency Processor User's Guide, product manual, Revision D 49

Tajima, O. 2008, Level-1 Data Format Table, https://cmb.uchicago.edu/project/images/e/e3/Level0_to_level1_format.pdf, QUIET internal memo 71

—. 2009a, Data Management at the Site, https://cmb.uchicago.edu/project/index.php/Data_Management_at_the_Site, QUIET internal memo 69

—. 2009b, Module temperature, e-mail 23

—. 2009c, Plan of Data flow and monitoring scheme for W-band, https://cmb.uchicago.edu/project/images/a/a0/Plan_of_dataflow_and_monitoring_for_Wband_20090517.pdf, QUIET internal memo 69

—. 2009d, Ps_curve-20090117.pdf, https://cmb.uchicago.edu/project/images/0/0b/Ps_curve-20090117.pdf, QUIET internal memo 37

—. 2010, W-band Sparse Wiregrid Analysis, https://cmb.uchicago.edu/project/images/3/3b/W91_wiregrid_analysis_2009.pdf, QUIET internal memo 93

—. 2011a, $I \rightarrow Q/U$ Monopole Leakage in W-band, https://cmb.uchicago.edu/project/images/2/24/2011_1121_leakage_monopole.pdf, QUIET internal memo 142, 150

—. 2011b, Systematic errors for multipole leakages, https://cmb.uchicago.edu/project/images/4/4a/2011_1220_multipole_leakages_tajimao.pdf, QUIET internal memo 142

—. 2012, Foreground Study Cross correlation QUIET x PSM-dusts, https://cmb.uchicago.edu/project/images/1/18/2012_0626_psm_x_quiet_pcl.pdf, QUIET internal memo 152

Tajima, O., Kusaka, A., Buder, I., Brizius, A., & Bischoff, C. 2010, How to complete the "QUIET Daily Checklist", https://cmb.uchicago.edu/site_document/trunk/checklist_how_to.pdf, QUIET internal memo 69

Tajima, O., Nguyen, H., Bischoff, C., et al. 2012, *Journal of Low Temperature Physics*, 167, 936–938, 81

Tucci, M., & Toffolatti, L. 2012, ArXiv e-prints, [arXiv:1204.0427 \[astro-ph.CO\]](https://arxiv.org/abs/1204.0427) 152

Vanderlinde, K. 2007, The QUIET DAQ Crate, https://cmb.uchicago.edu/project/images/6/64/QUIET_ADC.pdf, QUIET internal memo 53, 54, 60, 177

Weiland, J., Odegard, N., Hill, R., et al. 2011, *ApJS*, 192, 19 91

Zaldarriaga, M., & Seljak, U. 1997, *Phys. Rev. D*, 55, 1830 11